

US010026219B2

(12) **United States Patent**
Holzer et al.

(10) **Patent No.:** **US 10,026,219 B2**
(45) **Date of Patent:** **Jul. 17, 2018**

(54) **ANALYSIS AND MANIPULATION OF PANORAMIC SURROUND VIEWS**

(71) Applicant: **Fyusion, Inc.**, San Francisco, CA (US)

(72) Inventors: **Stefan Johannes Josef Holzer**, San Mateo, CA (US); **Stephen David Miller**, San Francisco, CA (US); **Radu Bogdan Rusu**, San Francisco, CA (US)

(73) Assignee: **FYUSION, INC.**, San Francisco, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **15/802,437**

(22) Filed: **Nov. 2, 2017**

(65) **Prior Publication Data**

US 2018/0068485 A1 Mar. 8, 2018

Related U.S. Application Data

(63) Continuation of application No. 14/539,889, filed on Nov. 12, 2014, now Pat. No. 9,836,873.

(Continued)

(51) **Int. Cl.**

G06T 15/00 (2011.01)

G06T 15/20 (2011.01)

(Continued)

(52) **U.S. Cl.**

CPC **G06T 15/205** (2013.01); **G06F 3/04842** (2013.01); **G06F 17/30256** (2013.01); **G06F 17/30277** (2013.01); **G06F 17/30837** (2013.01); **G06T 3/4038** (2013.01); **G06T 17/00** (2013.01); **H04N 5/23238** (2013.01); **H04N 5/265** (2013.01); **H04N 13/279** (2018.05);

(Continued)

(58) **Field of Classification Search**

CPC G06T 15/205; G06F 3/4038

USPC 348/36

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,746,376 B2 6/2010 Mendoza et al.

8,254,684 B2 8/2012 Raju

(Continued)

FOREIGN PATENT DOCUMENTS

CN 105849781 A 8/2016

DE 112014005165 T5 7/2016

(Continued)

OTHER PUBLICATIONS

“U.S. Appl. No. 14/530,669, Advisory Action dated Aug. 8, 2017”, 5 pgs.

(Continued)


Primary Examiner — Jeffery Williams

(74) *Attorney, Agent, or Firm* — Kwan & Olynick LLP

(57) **ABSTRACT**

Various embodiments of the present invention relate generally to systems and methods for analyzing and manipulating images and video. According to particular embodiments, the spatial relationship between multiple images and video is analyzed together with location information data, for purposes of creating a representation referred to herein as a surround view. In particular embodiments, a surround view can be generated by combining a panoramic view of an object with a panoramic view of a distant scene, such that the object panorama is placed in a foreground position relative to the distant scene panorama. Such combined panoramas can enhance the interactive and immersive viewing experience of the surround view.

20 Claims, 37 Drawing Sheets



Related U.S. Application Data

- (60) Provisional application No. 61/903,359, filed on Nov. 12, 2013.

- (51) **Int. Cl.**
 - G06T 17/00** (2006.01)
 - H04N 13/279** (2018.01)
 - H04N 13/282** (2018.01)
 - H04N 5/232** (2006.01)
 - G06F 17/30** (2006.01)
 - G06F 3/0484** (2013.01)
 - G06T 3/40** (2006.01)
 - H04N 5/265** (2006.01)

- (52) **U.S. Cl.**
 - CPC . **H04N 13/282** (2018.05); **G06T 2207/10021** (2013.01); **G06T 2207/20221** (2013.01); **G06T 2215/12** (2013.01); **G06T 2215/16** (2013.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,452,794 B2	5/2013	Yang et al.
8,705,892 B2	4/2014	Aguilera et al.
8,788,587 B2	7/2014	Balassanian et al.
2006/0188147 A1	8/2006	Rai et al.
2008/0201734 A1	8/2008	Lyon et al.
2008/0246759 A1	10/2008	Summers
2010/0002071 A1	1/2010	Ahiska
2010/0082665 A1	4/2010	Tanaka
2010/0171691 A1	7/2010	Cook et al.
2010/0191459 A1	7/2010	Carter et al.
2010/0208029 A1	8/2010	Marti et al.
2011/0098917 A1	4/2011	LeBeau et al.
2011/0110605 A1	5/2011	Cheong
2011/0184949 A1	7/2011	Luo
2011/0254835 A1	10/2011	Segal
2011/0316970 A1	12/2011	Cheong
2012/0019557 A1	1/2012	Aronsson et al.
2012/0293607 A1	11/2012	Bhogal et al.
2012/0307003 A1	12/2012	Huang et al.
2013/0046749 A1	2/2013	Bennett
2013/0050188 A1	2/2013	Fujiwara
2013/0155180 A1	6/2013	Wantland et al.
2013/0162634 A1	6/2013	Baik
2013/0250045 A1	9/2013	Ki et al.
2013/0346068 A1	12/2013	Solem et al.
2014/0046794 A1	2/2014	Vallery et al.
2014/0253436 A1	9/2014	Petersen
2014/0362274 A1	12/2014	Christie et al.
2015/0019643 A1	1/2015	Balassanian et al.
2015/0036945 A1	2/2015	Zuliani
2015/0039616 A1	2/2015	Rolston et al.
2015/0130799 A1	5/2015	Holzer et al.
2015/0130800 A1	5/2015	Holzer et al.
2015/0130894 A1	5/2015	Holzer et al.
2015/0134651 A1	5/2015	Holzer et al.
2015/0138190 A1	5/2015	Holzer et al.
2015/0153933 A1	6/2015	Filip et al.
2015/0339846 A1	11/2015	Holzer et al.
2016/0005435 A1	1/2016	Campbell et al.

FOREIGN PATENT DOCUMENTS

GB	2534821 A	8/2016
WO	2015073570 A2	5/2015

OTHER PUBLICATIONS

- “U.S. Appl. No. 14/530,669, Examiner Interview Summary dated Apr. 14, 2017”, 3 pages.
- “U.S. Appl. No. 14/530,669, Examiner Interview Summary dated Aug. 8, 2017”, 2 pgs.

- “U.S. Appl. No. 14/530,669, Final Office Action dated Apr. 20, 2017”, 25 pages.
- “U.S. Appl. No. 14/530,669, Non Final Office Action dated Jan. 3, 2017”, 26 pages.
- “U.S. Appl. No. 14/530,671, Advisory Action dated Aug. 4, 2017”, 7 pgs.
- “U.S. Appl. No. 14/530,671, Examiner Interview Summary dated Apr. 14, 2017”, 3 pages.
- “U.S. Appl. No. 14/530,671, Final Office Action dated Apr. 20, 2017”, 31 pages.
- “U.S. Appl. No. 14/530,671, Non Final Office Action dated Jan. 3, 2017”, 32 pages.
- “U.S. Appl. No. 14/530,673, Final Office Action dated Oct. 4, 2017”, 8 pgs.
- “U.S. Appl. No. 14/530,673, Non Final Office Action dated Mar. 23, 2017”, 12 pages.
- “U.S. Appl. No. 14/539,814, Advisory Action dated Aug. 4, 2017”, 8 pgs.
- “U.S. Appl. No. 14/539,814, Examiner Interview Summary dated Apr. 17, 2017”, 3 pages.
- “U.S. Appl. No. 14/539,814, Final Office Action dated Apr. 20, 2017”, 38 pages.
- “U.S. Appl. No. 14/539,814, Non Final Office Action dated Dec. 30, 2016”, 37 pages.
- “U.S. Appl. No. 14/539,889, Advisory Action dated Jul. 19, 2017”, 3 pages.
- “U.S. Appl. No. 14/539,889, Final Office Action dated Apr. 13, 2017”, 12 pages.
- “U.S. Appl. No. 14/539,889, Non Final Office Action dated Oct. 6, 2016”, 14 pages.
- “U.S. Appl. No. 14/539,889, Notice of Allowance dated Aug. 3, 2017”, 9 pgs.
- “U.S. Appl. No. 14/819,473, Advisory Action dated May 30, 2017”, 7 pgs.
- “U.S. Appl. No. 14/819,473, Examiner Interview Summary dated Apr. 18, 2017”, 3 pages.
- “U.S. Appl. No. 14/819,473, Examiner Interview Summary dated Jul. 7, 2017”, 3 pages.
- “U.S. Appl. No. 14/819,473, Examiner Interview Summary dated Jul. 11, 2016”, 3 pages.
- “U.S. Appl. No. 14/819,473, Examiner Interview Summary dated Aug. 17, 2016”, 3 pages.
- “U.S. Appl. No. 14/819,473, Examiner Interview Summary dated Oct. 14, 2016”, 3 pages.
- “U.S. Appl. No. 14/819,473, Final Office Action dated Mar. 10, 2017”, 15 pages.
- “U.S. Appl. No. 14/819,473, Final Office Action dated Apr. 28, 2016”, 45 pages.
- “U.S. Appl. No. 14/819,473, Non Final Office Action dated Aug. 10, 2017”, 44 pgs.
- “U.S. Appl. No. 14/819,473, Non Final Office Action dated Sep. 1, 2016”, 36 pages.
- “U.S. Appl. No. 14/819,473, Non Final Office Action dated Oct. 8, 2015”, 44 pgs.
- “iFotosoft, “Tricks for Shooting iPhone Panoramas””, <https://www.iffotosoft.com>, Retrieved from internet: Mar. 20, 2017, pp. 1-7.
- “Int’l Application Serial No. PCT/US2014/065282, Search Report & Written Opinion dated Feb. 23, 2015”, 9 pgs.
- “Seene, “Introducing Seene””, <https://www.youtube.com/watch?v=JADH8PICxHc>, Uploaded Oct. 9, 2013, app demonstration video.
- “Seene, “Seene—Oe—Realtime Augmented Reality Technology””, Uploaded Feb. 1, 2012, app demonstration video, retrieved from <https://www.youtube.com/watch?v=UNiyT7qwHGA>.
- “The Octilus, “Visor Scubo 3D—Scubo 3D Viewer. Tutorial””, Uploaded Aug. 28, 2012, app tutorial video, retrieved from https://www.youtube.com/watch?v=J_bsfvYwic.
- “United Kingdom Application Serial No. 1609577.0, Office Action dated Jun. 7, 2017”, 2 pgs.
- “United Kingdom Application Serial No. 1609577.0, Office Action dated Jun. 15, 166”, 1 pg.
- Ballan, Luca et al., “Unstructured Video-Based Rendering: Interactive Exploration of Casually Captured Videos”, ACM, ACM

(56)

References Cited

OTHER PUBLICATIONS

Transactions on Graphics (TOG)—Proceedings of ACM SIGGRAPH 2010, vol. 29, Issue 4, Article No. 87, Jul. 30, 2010.

Cao, Xun et al., “Semi-Automatic 2D-to-3D Conversion Using Disparity Propagation”, IEEE, IEEE Transactions on Broadcasting, vol. 57, Issue 2, Apr. 19, 2011, pp. 491-499.

Chan, Shing-Chow et al., “An Object-Based Approach to Image/Video-Based Synthesis and Processing for 3-D and Multiview Televisions”, IEEE, IEEE Transactions on Circuits and Systems for Video Technology, vol. 19, Issue 6, Mar 16, 2009, pp. 821-831.

Chen, Shenchang E. , “QuickTime VR—An Image-Based Approach to Virtual Environment Navigation”, ACM, SIGGRAPH '95 Proceedings of the 22nd annual conference on Computer graphics and interactive techniques, 1995, 29-38.

Clare, Adam , “Reality is a Game; What is Skybox?”, retrieved from the Internet <<http://www.realityisagame.com/archives/1776/what-is-a-skybox/>>, 5.

Cole, David M. et al., “Using Laser Range Data for 3D SLAM in Outdoor Environments”, Proceedings of the IEEE International Conference on Robotics and Automation, May 2006, pp. 1556-1563.

Figueroa, Nadia et al., “From Sense to Print: Towards Automatic 3D Printing from 3D Sensing Devices”, IEEE, 2013 IEEE International Conference on Systems, Man, and Cybernetics (SMC, Oct. 13, 2013.

Golovinsky, Aleksey et al., “Shape-based Recognition of 3D Point Clouds in Urban Environment”, IEEE, IEEE 12th International Conference on Computer Vision (ICCV), 2009, 2154-2161.

Gurdan, Tobias et al., “Spatial and Temporal Interpolation of Multi-View Image Sequences”, Department of Computer Science, Technische Universität München Ascending Technologies GmbH, Krailing, Germany, Section 2.3, Image Warping and Blending; Retrieved from the Internet <https://vision.in.tum.de/_media/spezial/bib/gurdan-et-al-gcpr-2014.pdf>, 12.

Keller, Maik et al., “Real-Time 3D Reconstruction in Dynamic Scenes Using Point-Based Fusion”, IEEE, 2013 International Conference on 3DTV, Jul. 1, 2013.

Kim, Bo Y. , “Adaptive Strip Compression for Panorama Video Streaming”, Proceedings of the Computer Graphics International, 2004, 6 pgs.

Kottamasu, V. L. P. , “User Interaction of One-Dimensional Panoramic Images for iPod Touch”, Thesis, Linkoping University Electronic Press, LIU-IDA-LITH-EX-A—12/071—SE, Dec 4, 2012.

Mian, Ajmal S. et al., “Three-Dimensional Model-Based Object Recognition and Segmentation in Cluttered Scenes”, IEEE Transactions on Pattern Analysis and Machine Intelligence, col. 28, No. 10, Oct. 2006, 1584-1601.

Saito, Hideo et al., “View Interpolation of Multiple Cameras Based on Projective Geometry”, Department of Information and Computer Science, Keio University and Presto, Japan Science and Technology Corporation (JST), retrieved from the Internet <<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.6.5162&rep=rep1&type=pdf>>, 6.

Schiller, Ingo et al., “Datastructure for Capturing Dynamic Scenes with a Time-of-Flight Camera”, Springer-Verlad, Dyna3D 2009, LNCS 5742, 2009, 42-57.

Seitz, Steven M. , “Image-Based Transformation of Viewpoint and Scene Appearance”, A Dissertation submitted in partial fulfillment of the requirements for the Degree of Doctor of Philosophy at the University of Wisconsin; retrieved from the Internet <<http://homes.cs.washington.edu/~seitz/papers/thesis.pdf>>, 1997, 111.

Shade, Jonathan et al., “Layered Depth Images”, Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, ACM, SIGGRAPH, Jul. 24, 1998, pp. 231-242.

Snavely, Noah et al., “Photo Tourism: Exploring Phot Collections in 3D”, ACM, ACM Transactions on Graphics (TOG)—Proceeding of ACM SIGGRAPH 2006, vol. 25, Issue 3, Jul. 2006, 835-846.

Steder, Bastian et al., “Robust On-line Model-based Object Detection from Range Images”, International Conference on Intelligent Robots and Systems, pp. 4739-4744, Oct. 15, 2009, 6 pages.

Stella, Rick , “Become the Hero of Your GoPro With These Easy Tips and Tricks”, GoPro: 8 Tips and Tricks Digital Trends, Jul. 26, 2016, pp. 1-6.

Thyssen, Anthony , “ImageMagick v6 Examples—Color Basics and Channels”, Website http://www.imagemagick.org/Usage/color_basics/, Retrieved Dec. 23, 2016, Mar. 9, 2011.

Tompkin, James et al., “Videoscapes: Exploring Sparse, Unstructured Video Collections”, Acm Trans. Graph. 31 (4), Article 68, Jul. 2012, 12 pgs.

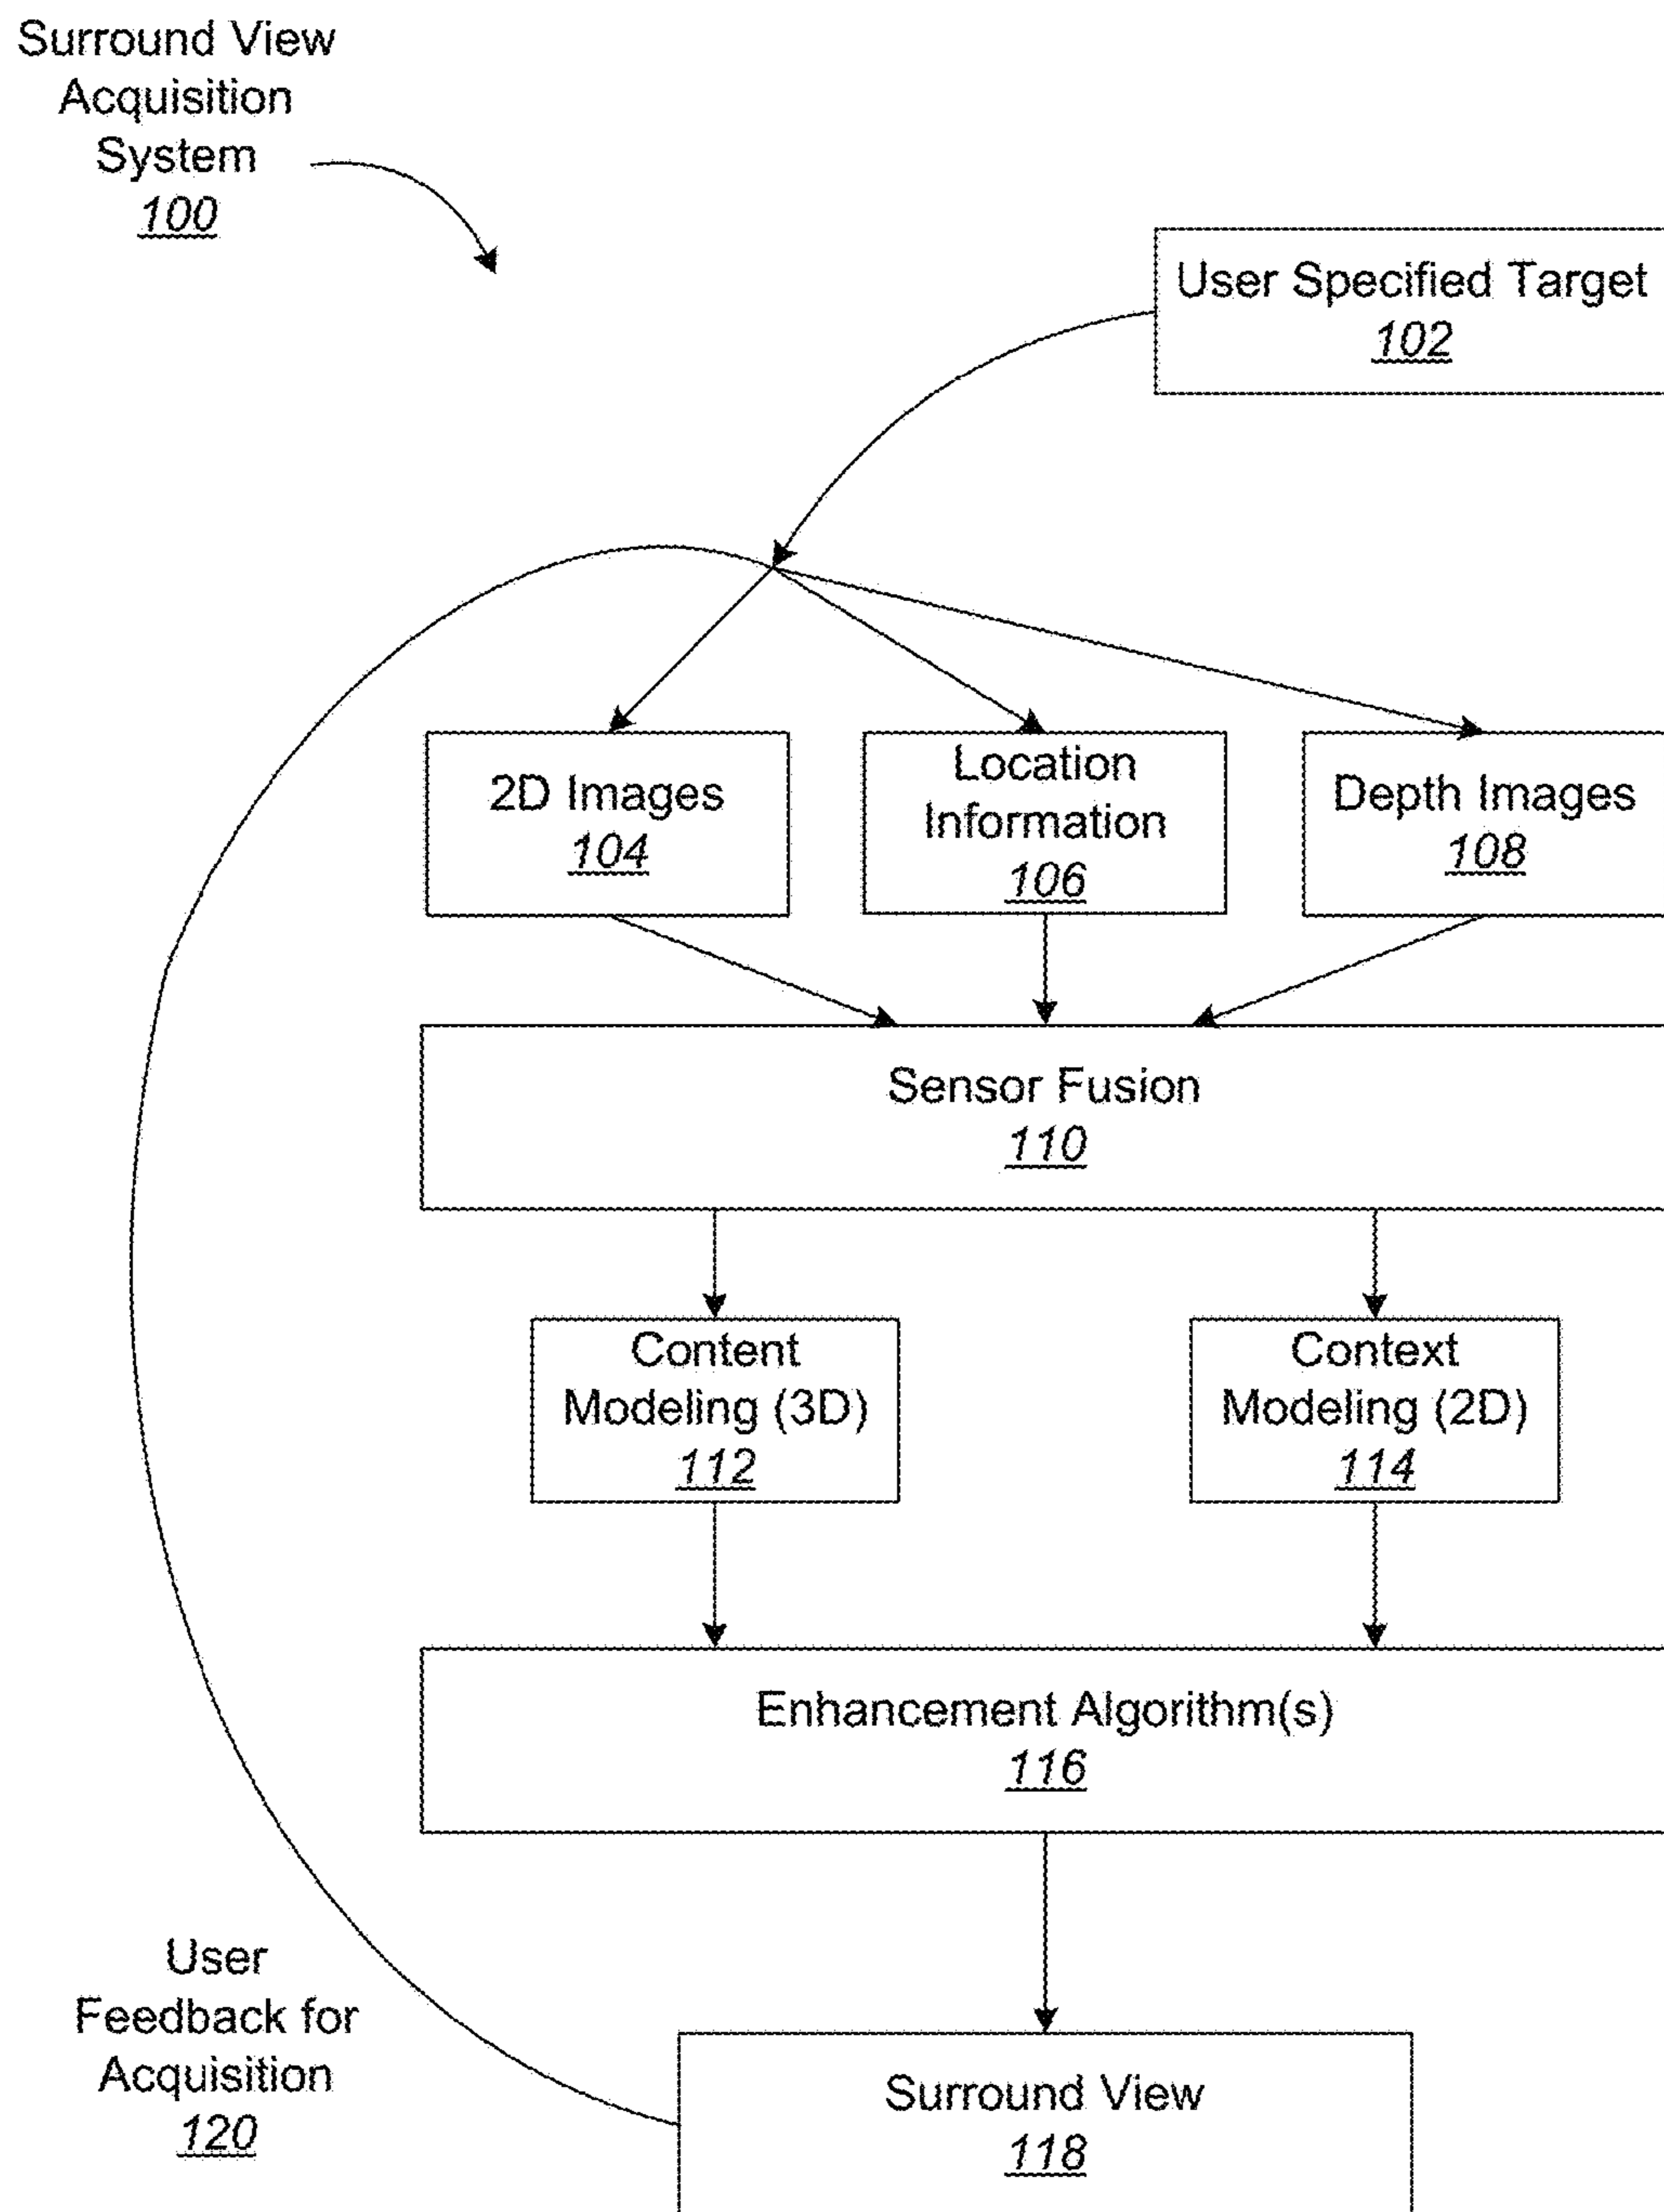
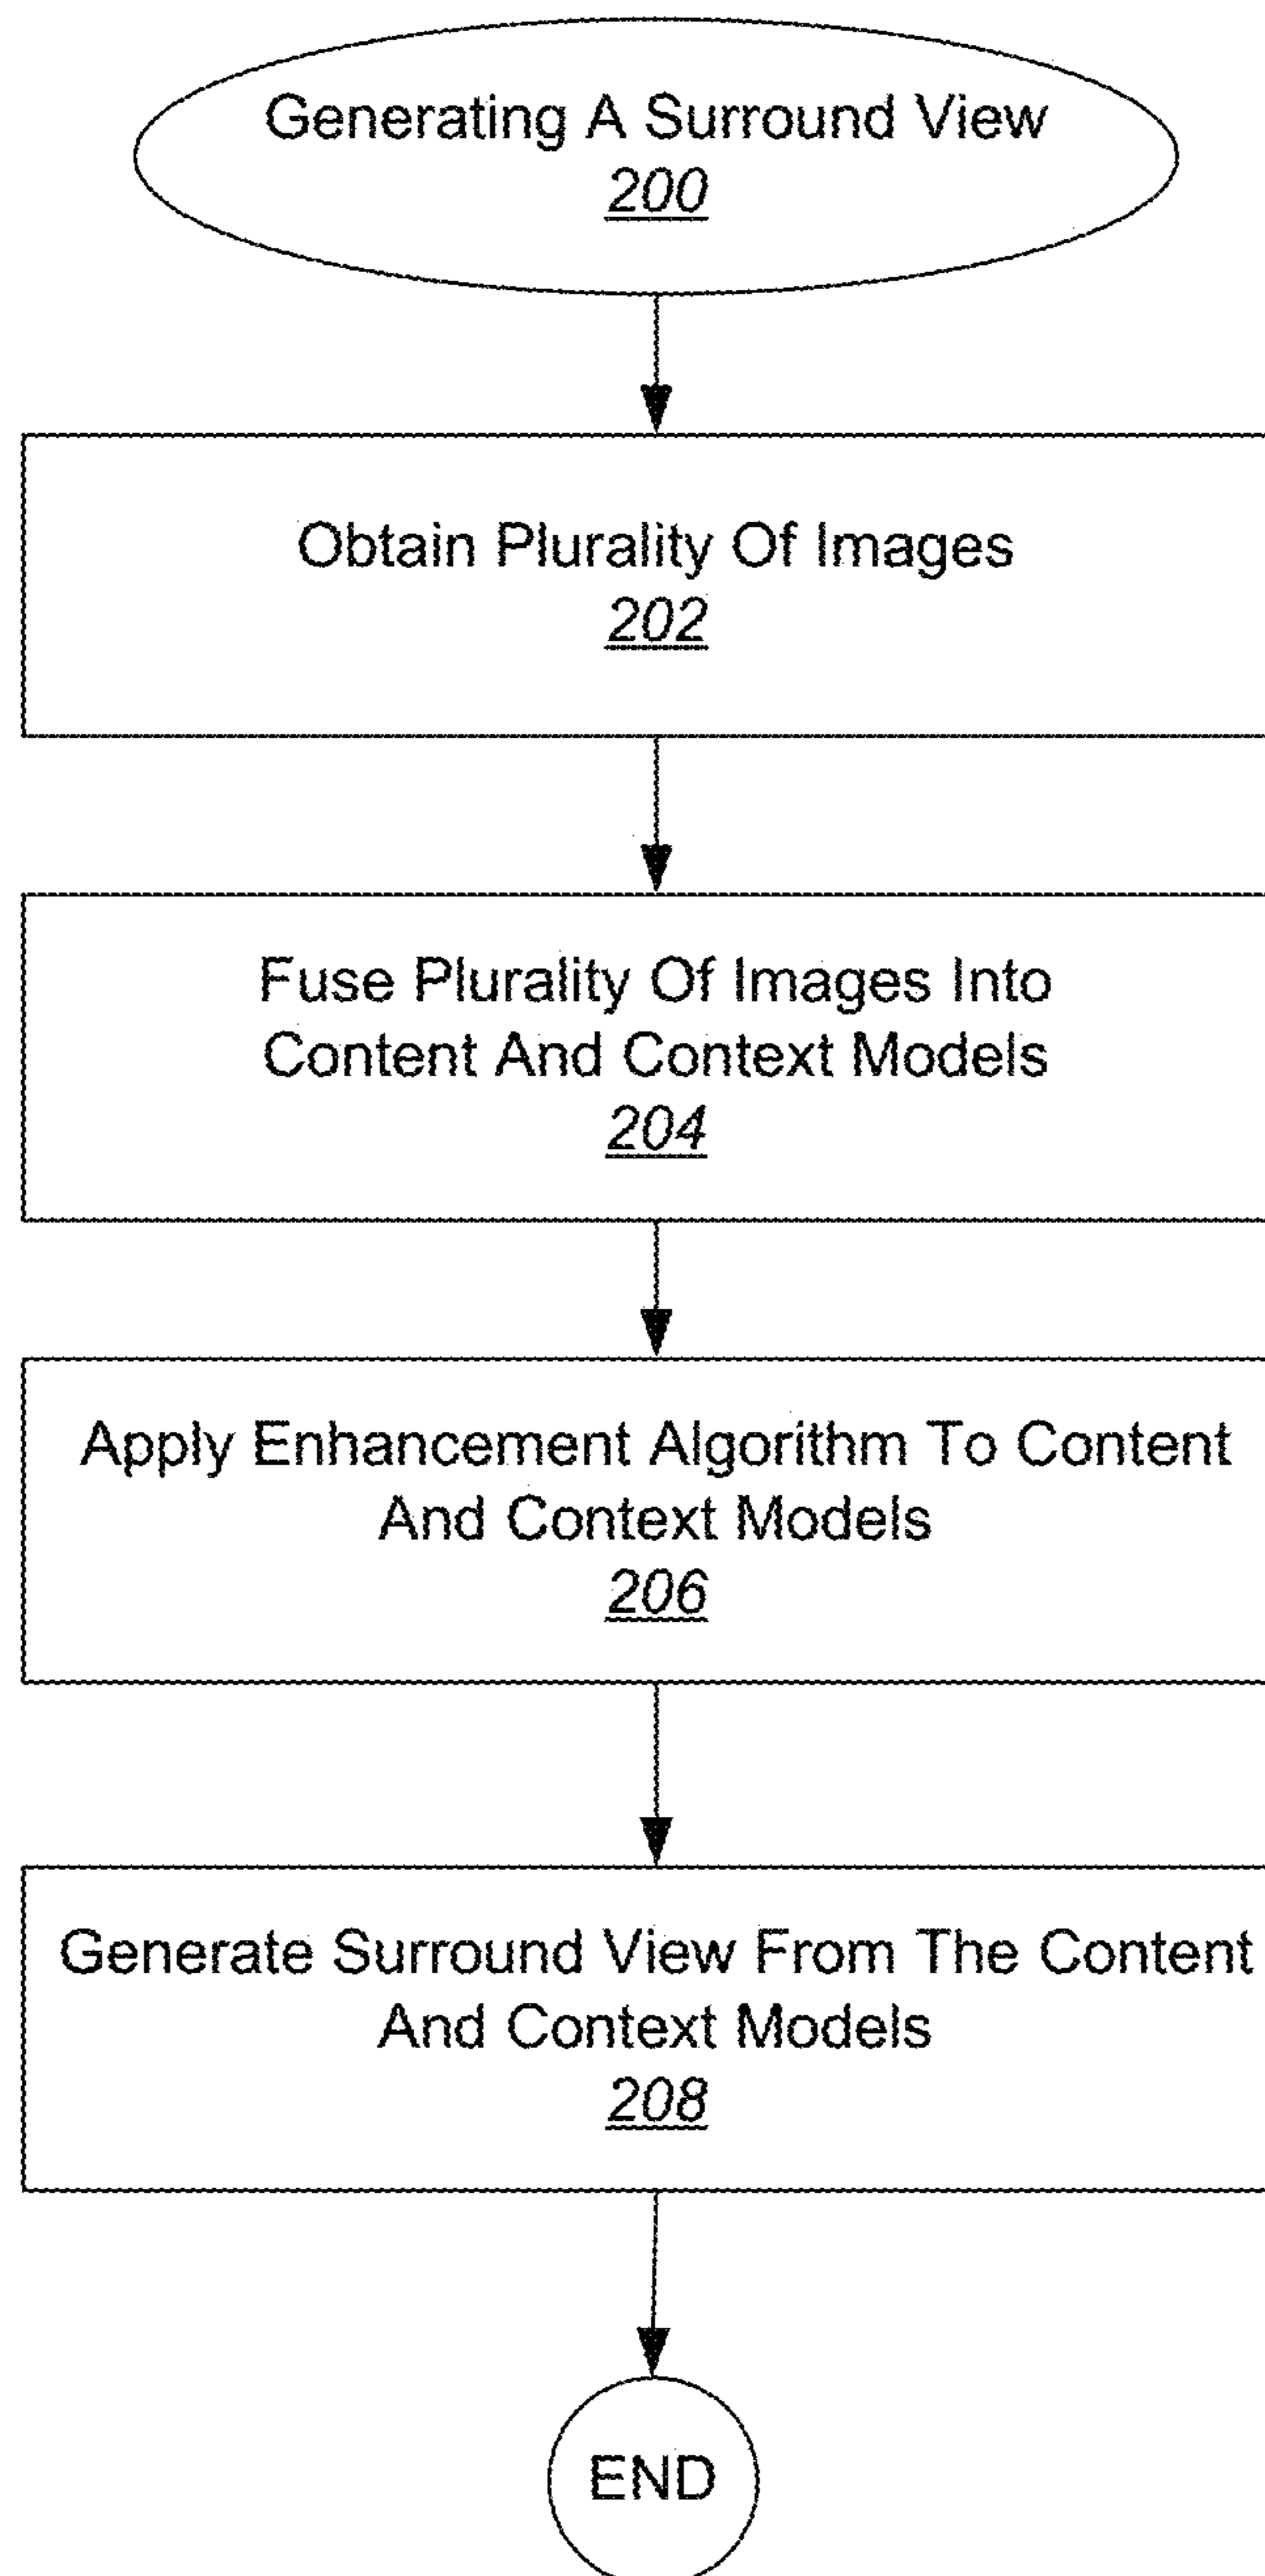


FIG. 1

**FIG. 2**




FIG. 3




FIG. 4A




FIG. 4B

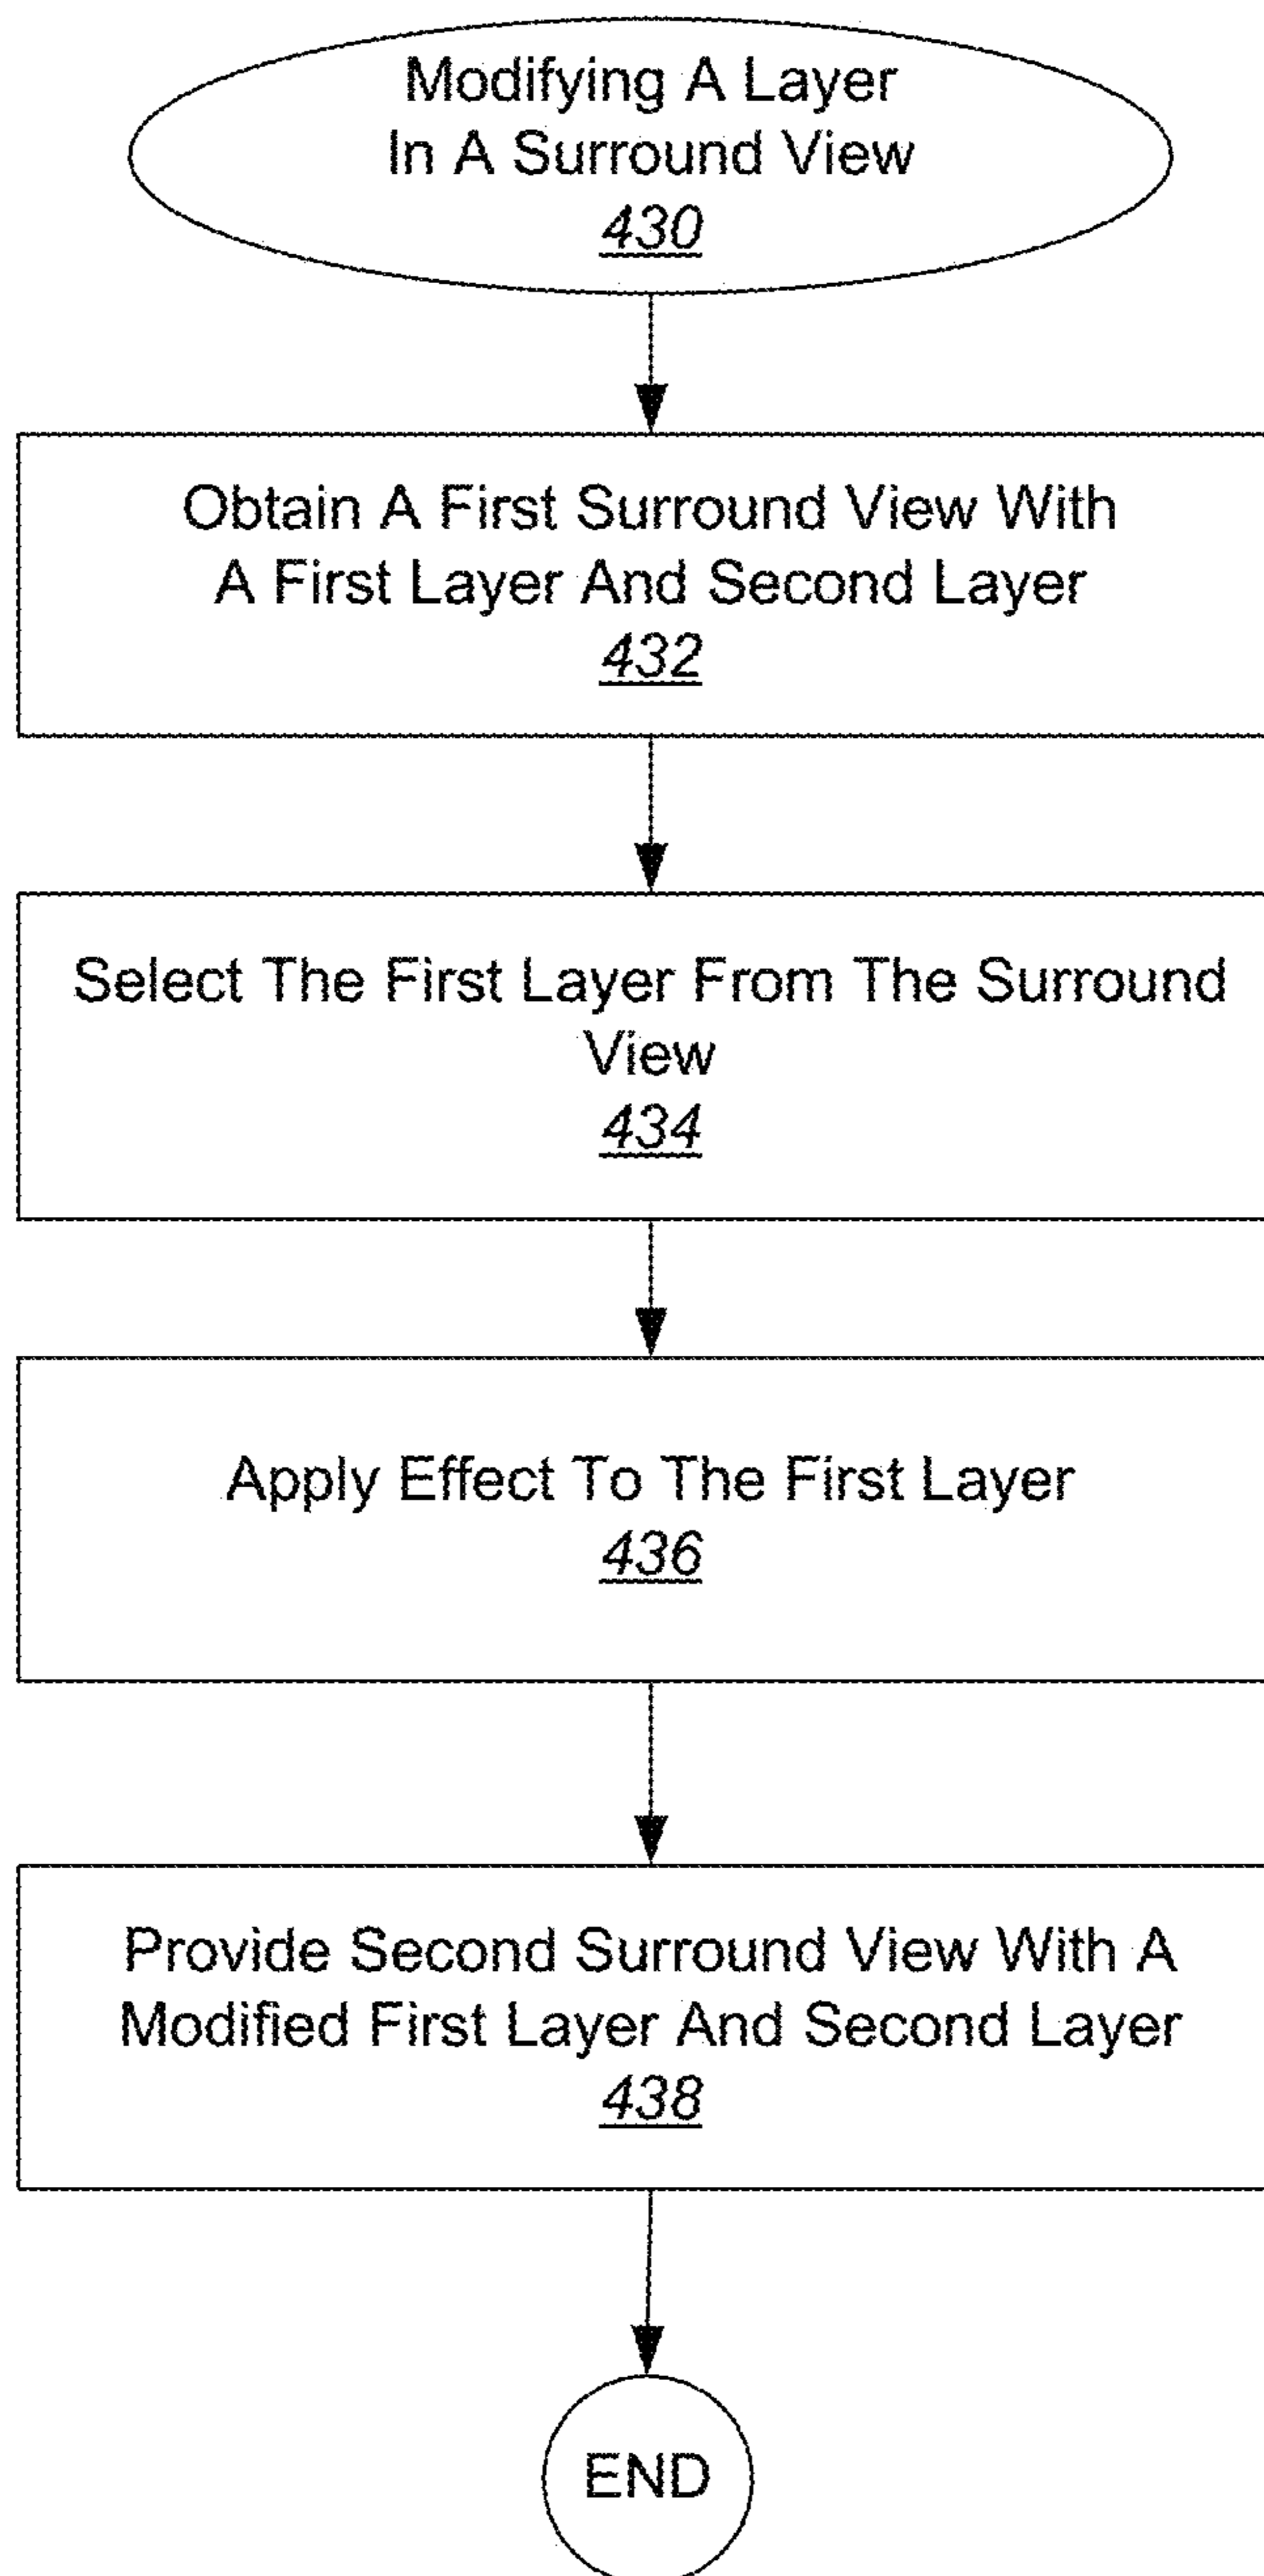


FIG. 4C




FIG. 5B

FIG. 5A




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 6D




FIG. 6E




FIG. 7A




FIG. 7B




FIG. 7C

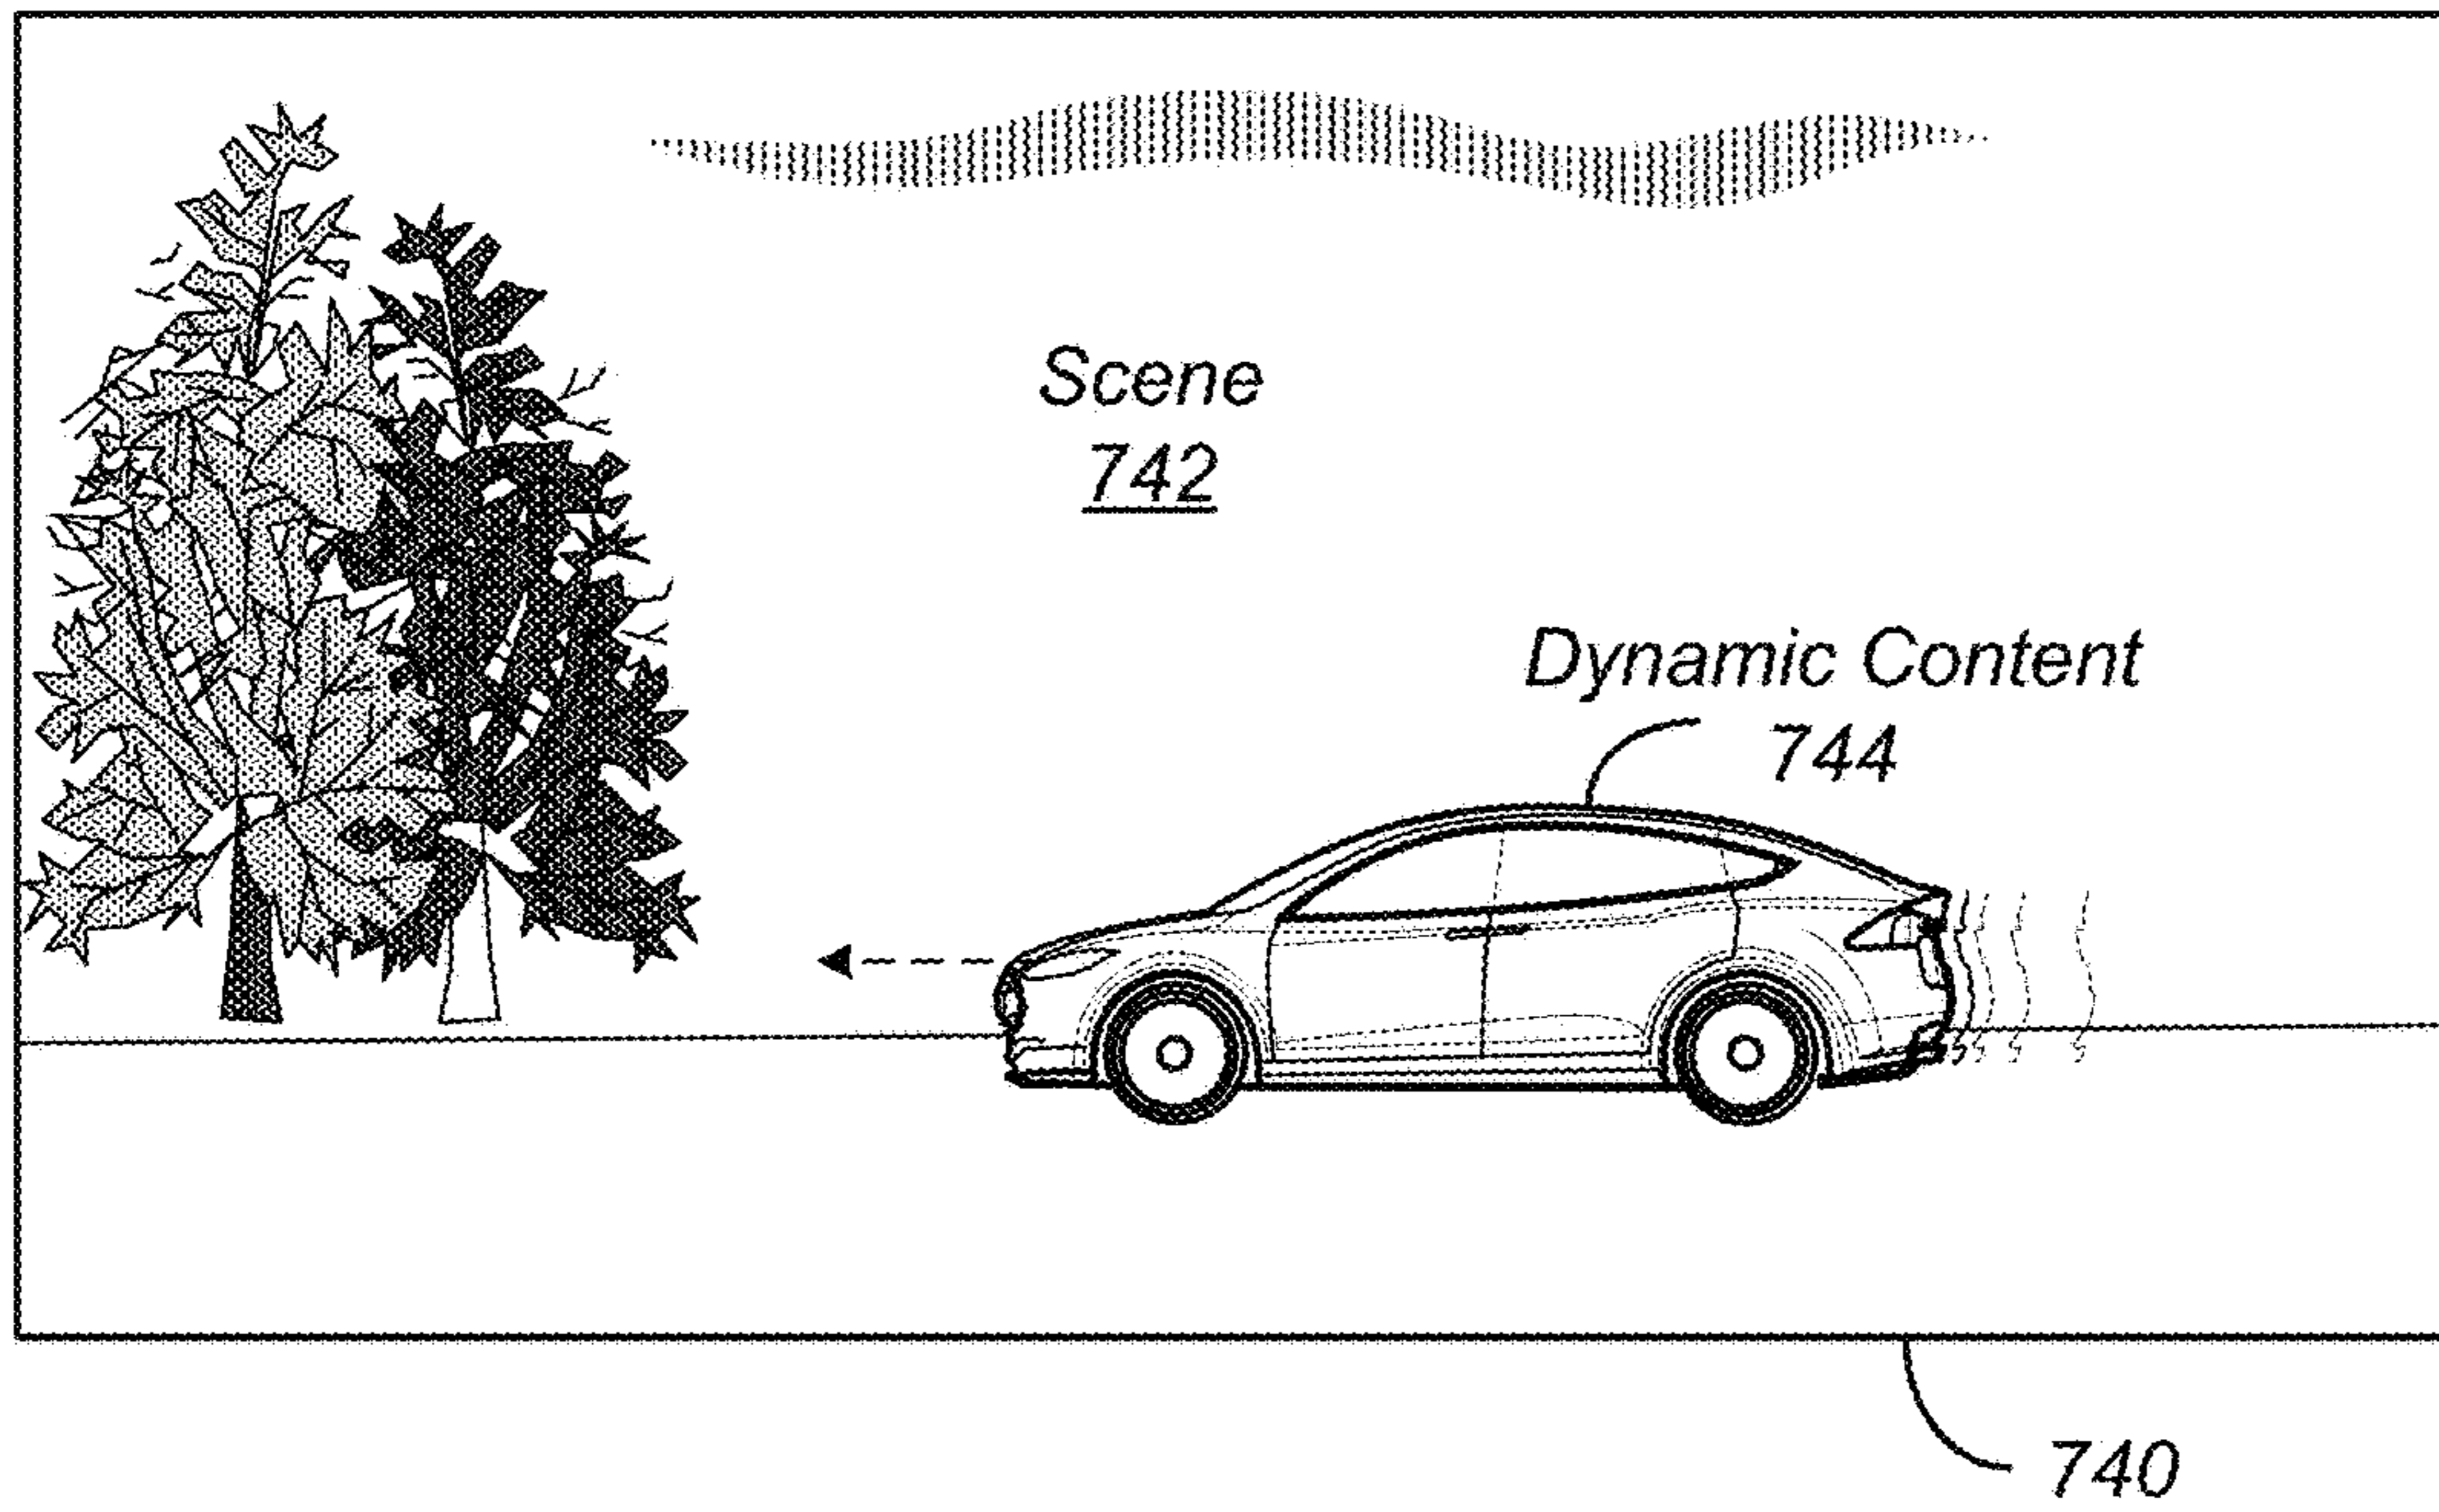


FIG. 7D




FIG. 7E




FIG. 7F




FIG. 7G




FIG. 7H

768




FIG. 7I

782

780

776

778




FIG. 7J




FIG. 7K




FIG. 7L




FIG. 7M




FIG. 8




FIG. 9




FIG. 10




FIG. 11

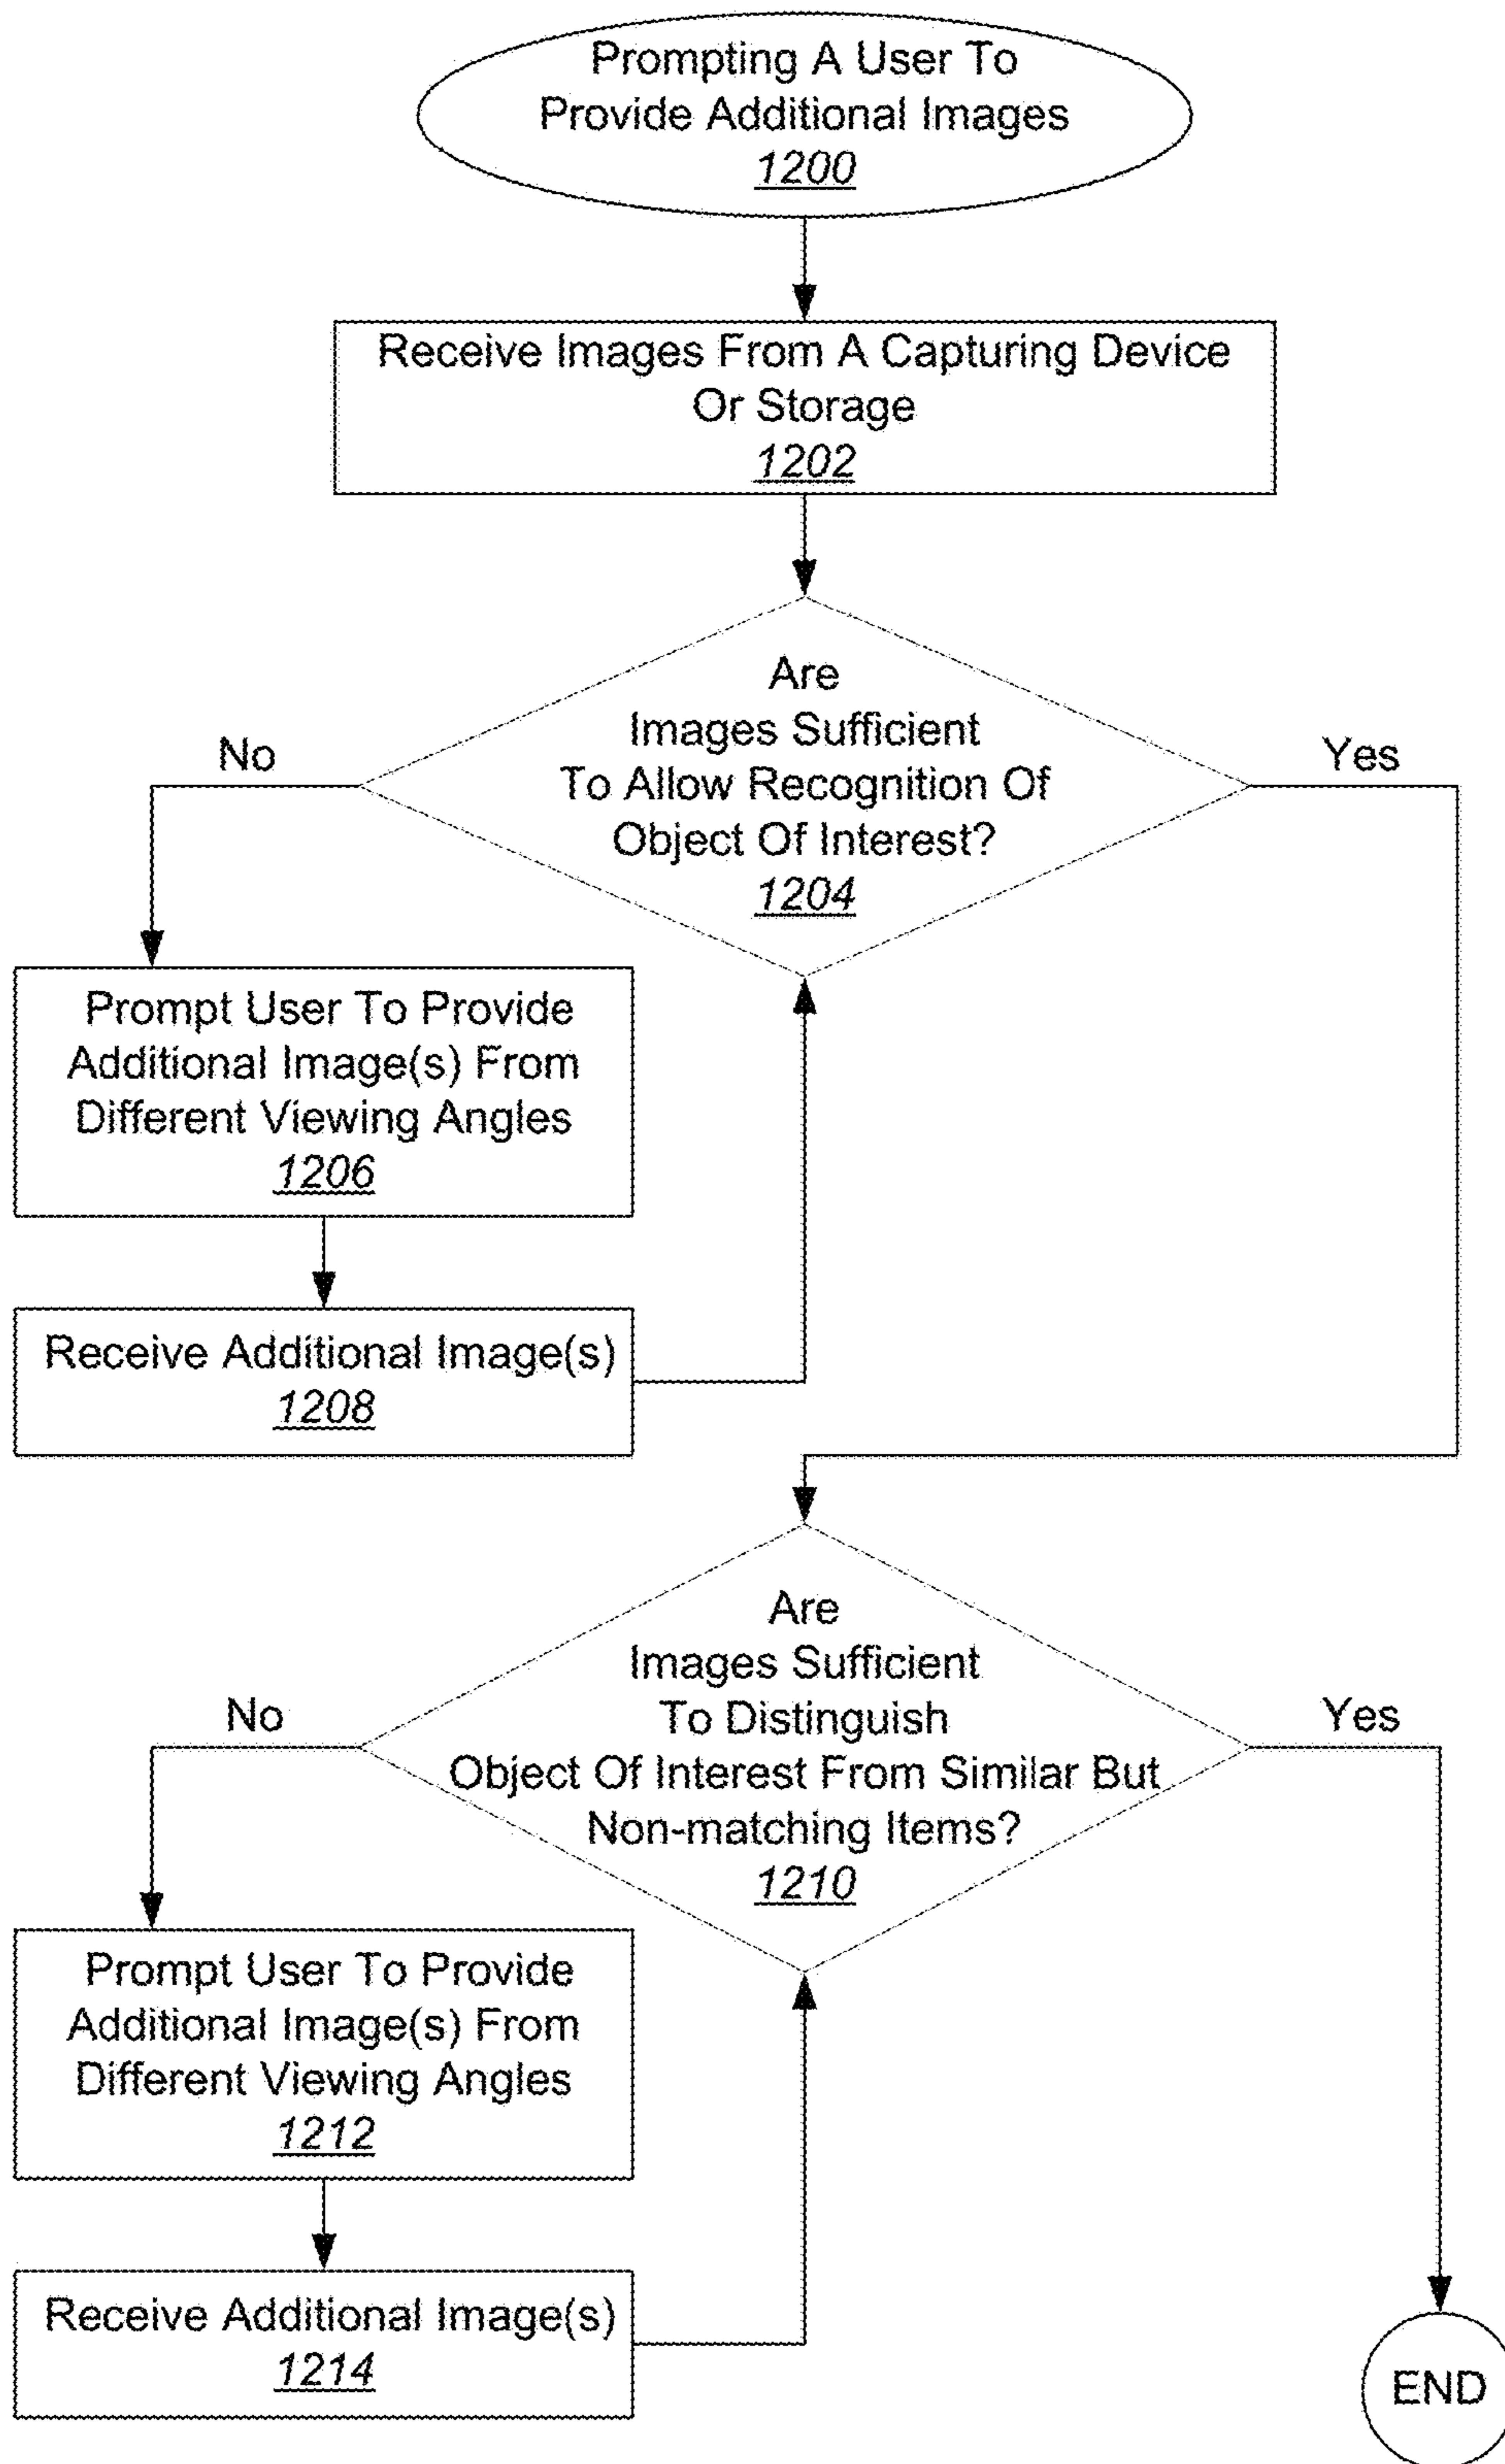


FIG. 12




FIG. 13A




FIG. 13B

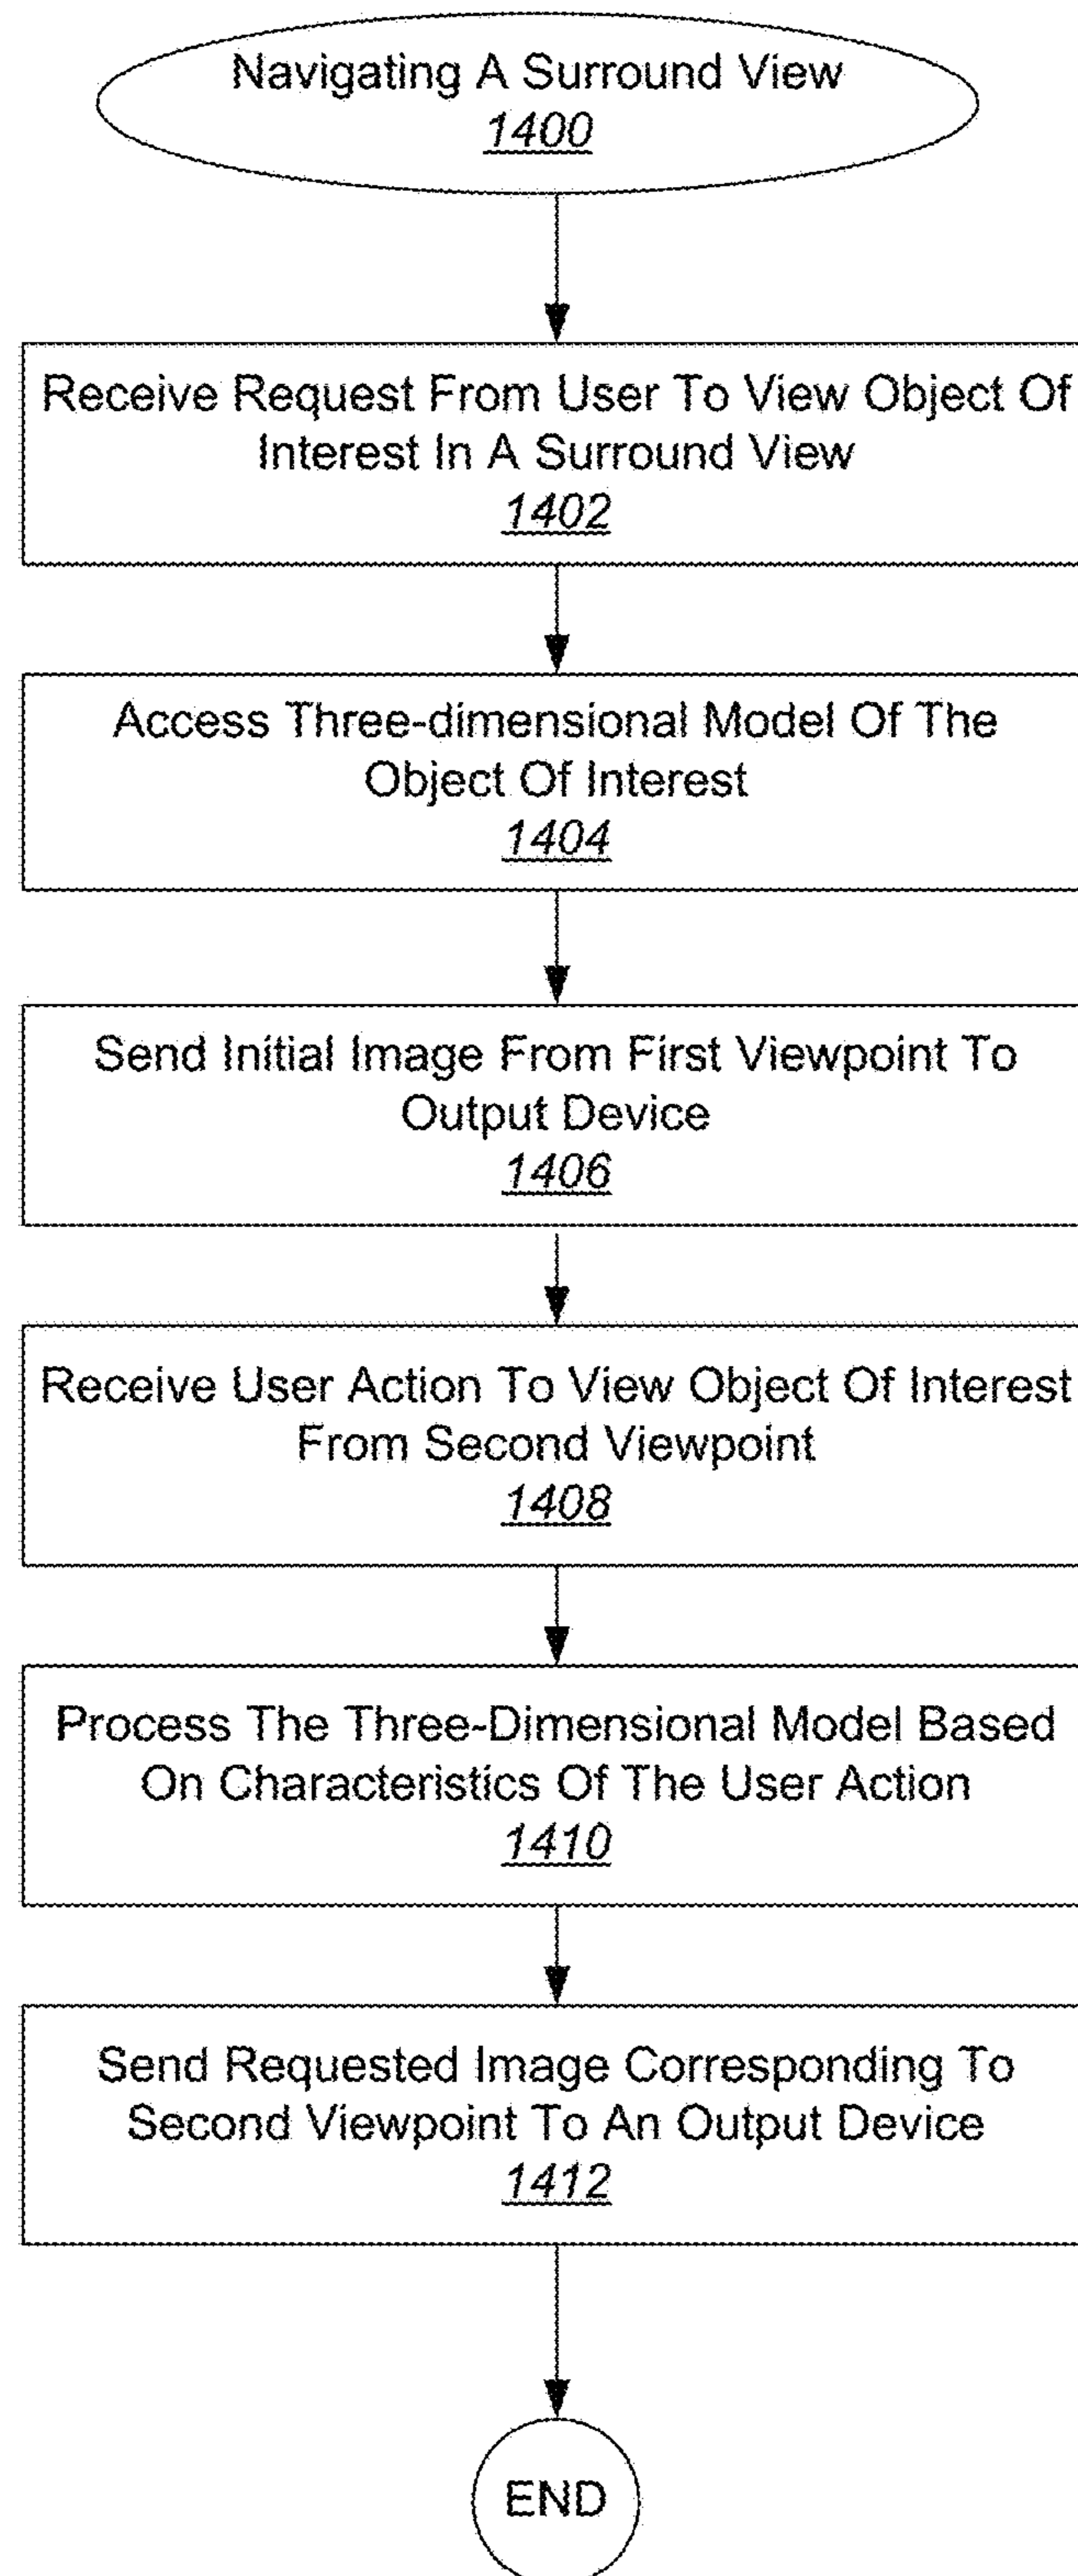


FIG. 14




FIG. 15




FIG. 16A




FIG. 16B

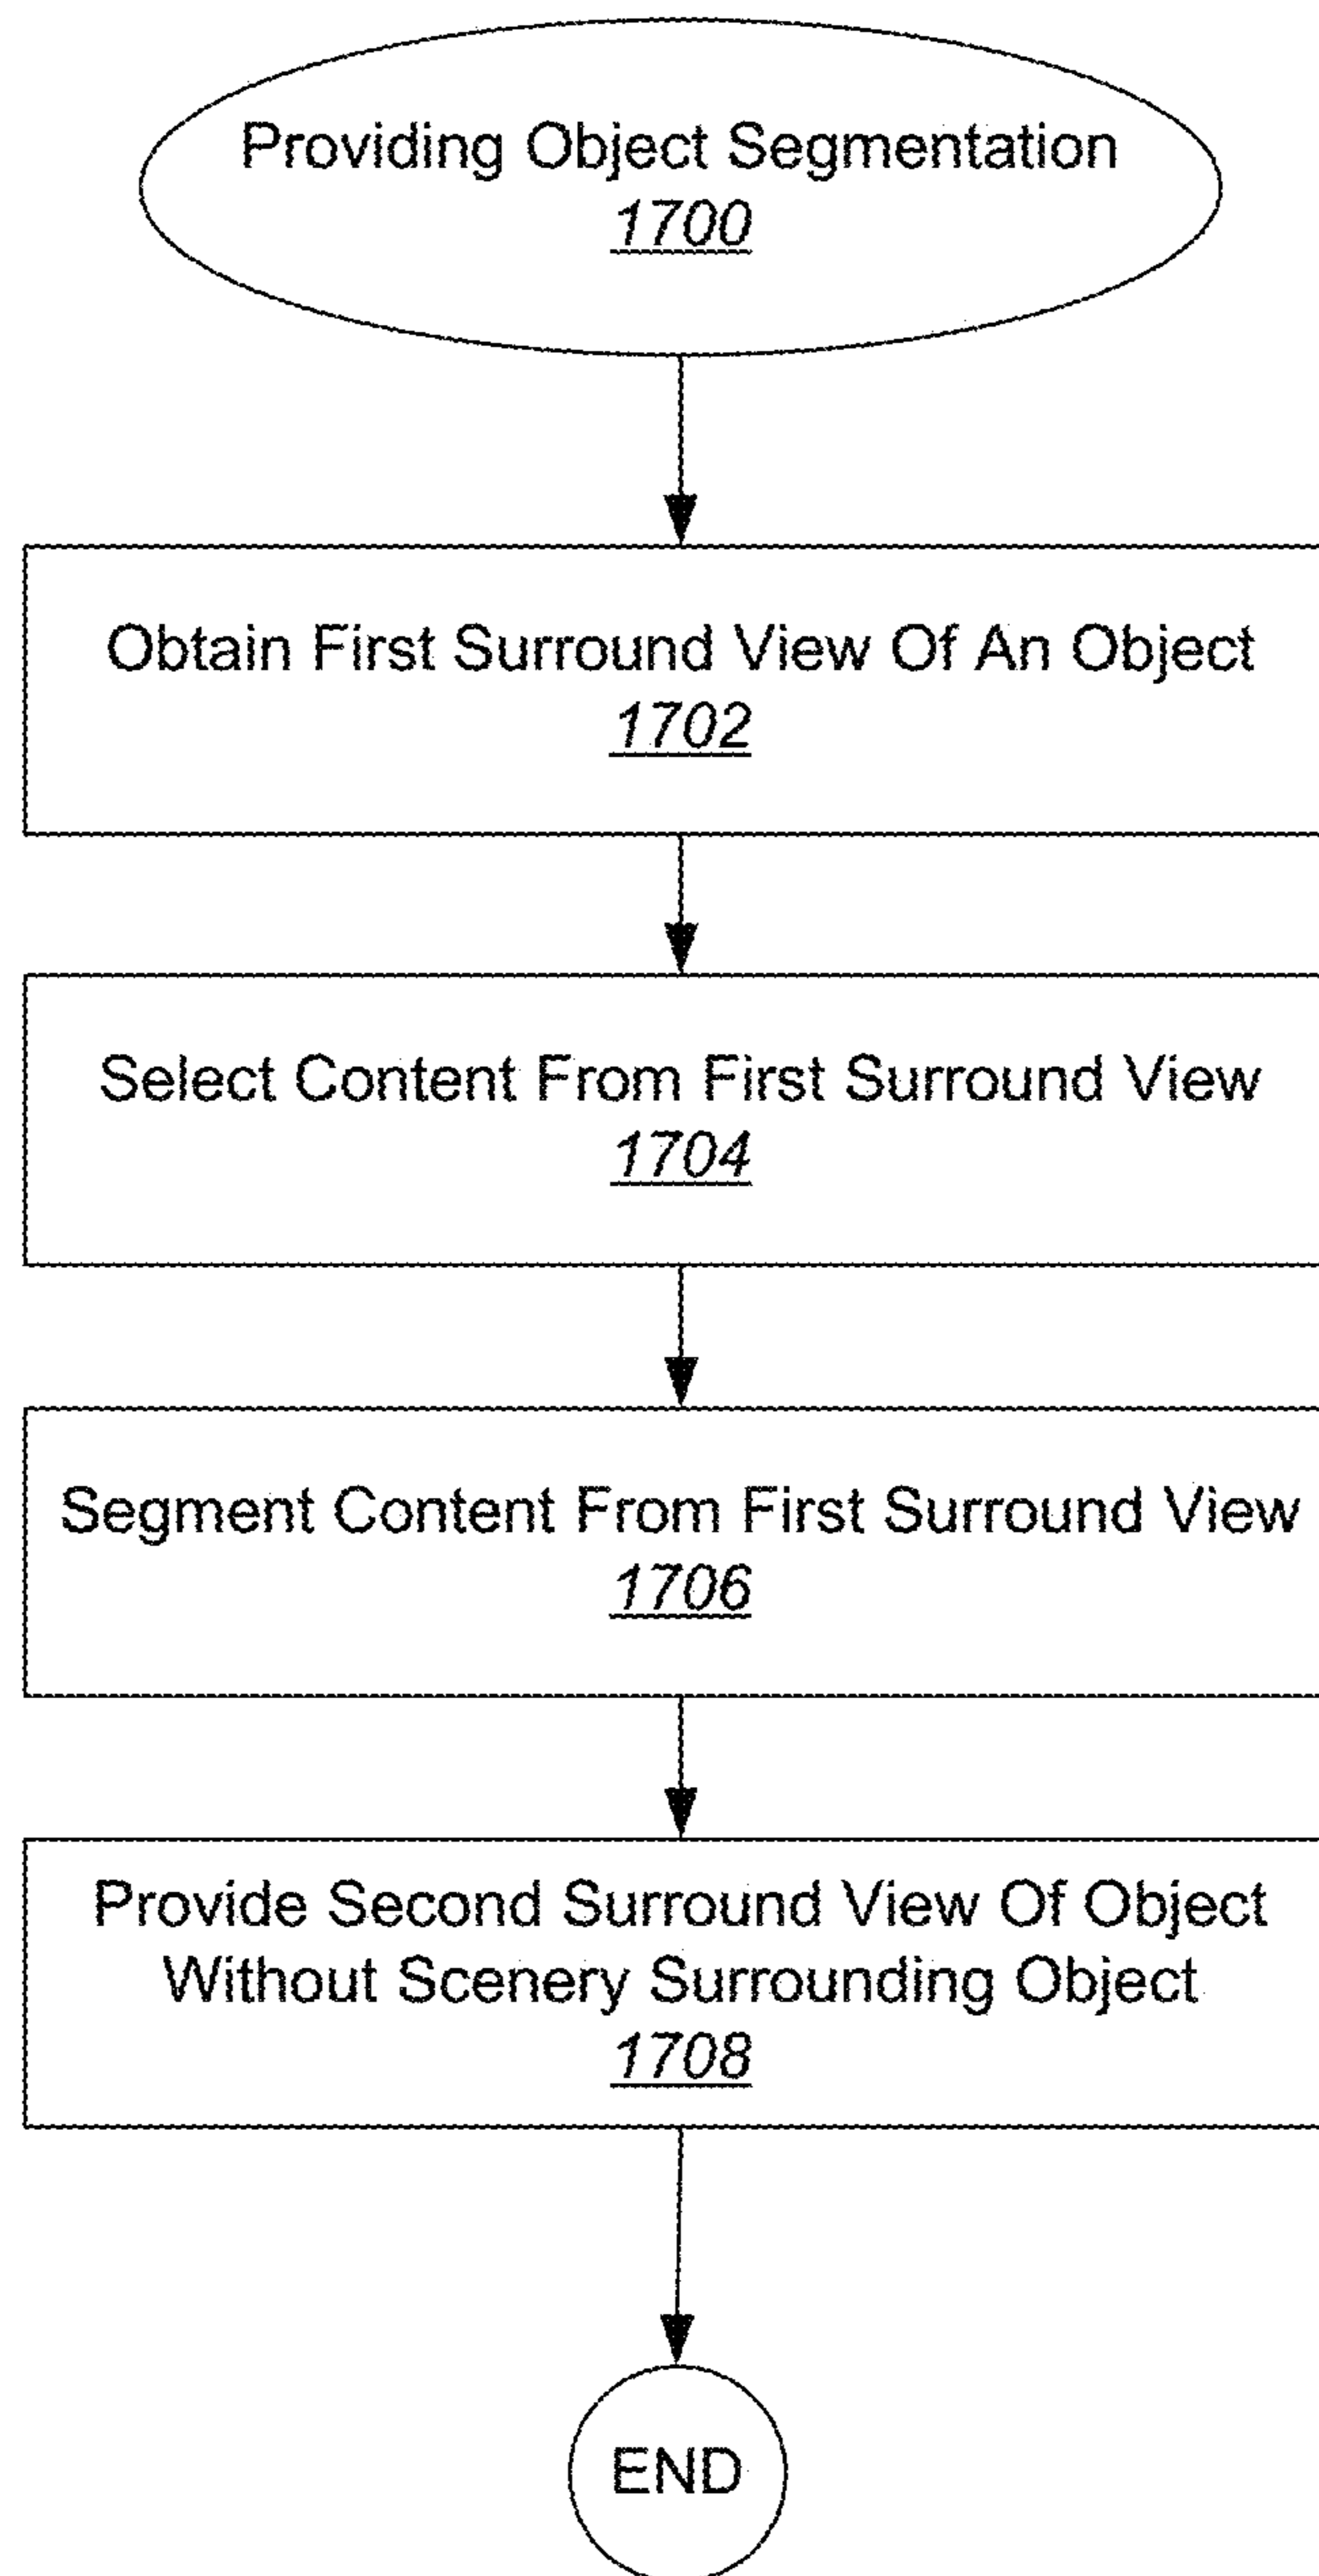


FIG. 17A




FIG. 17B




FIG. 18

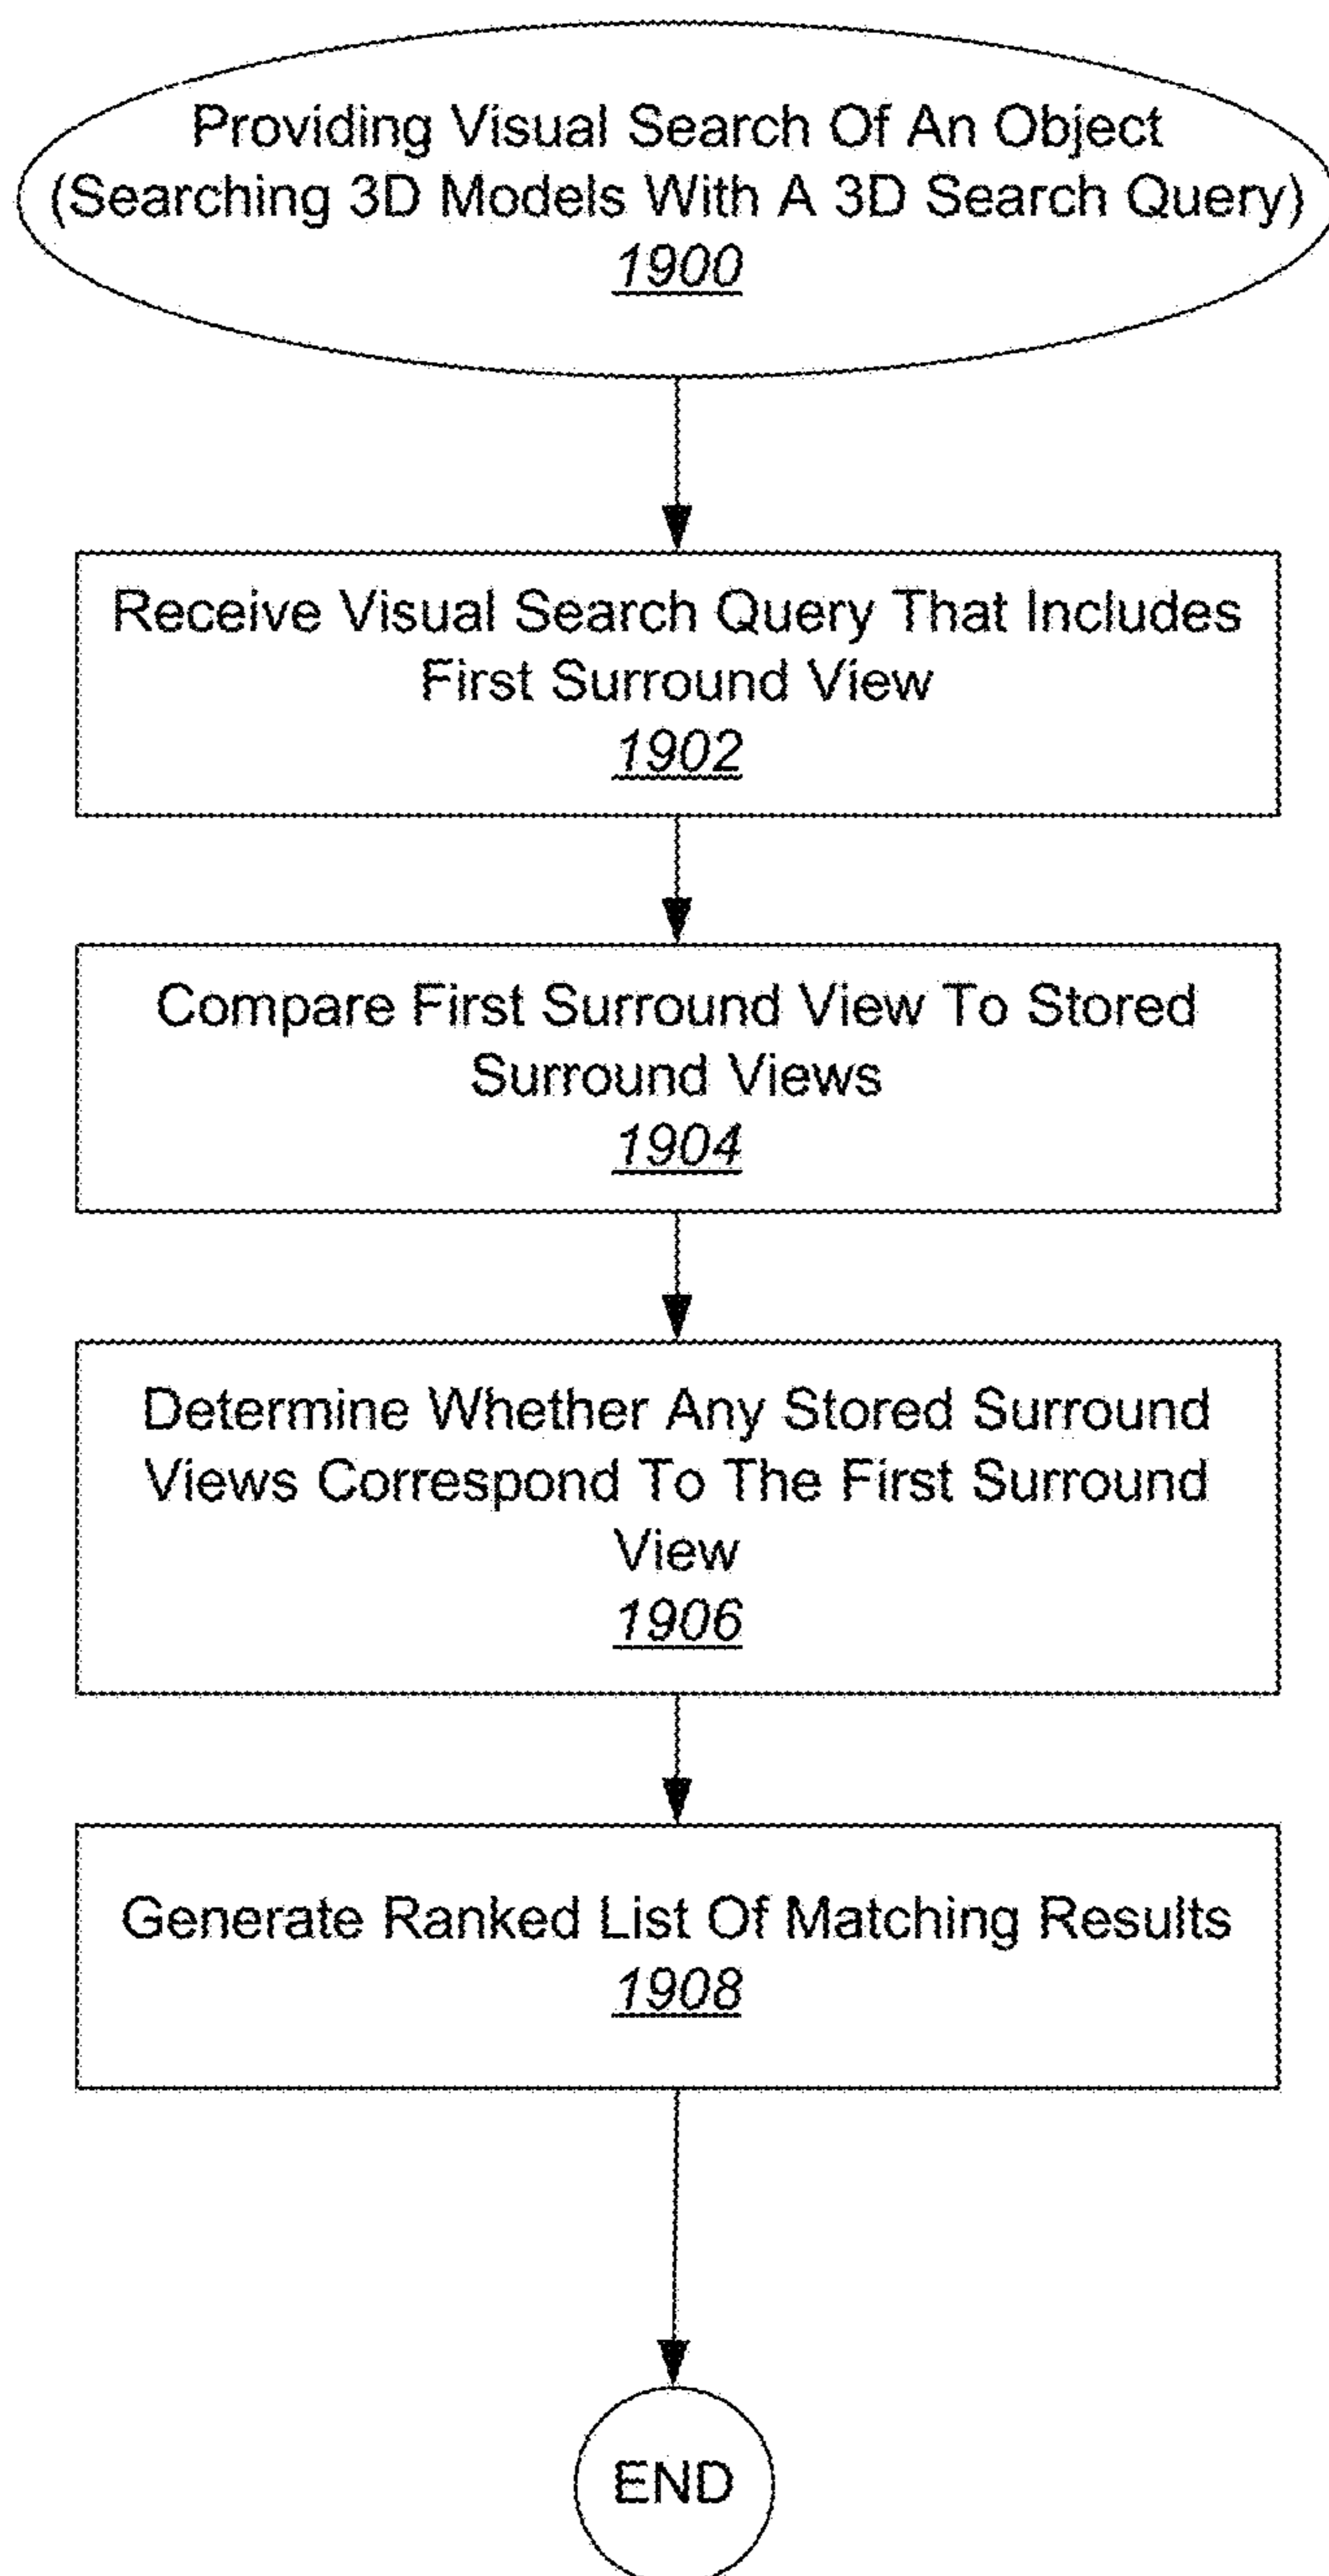


FIG. 19

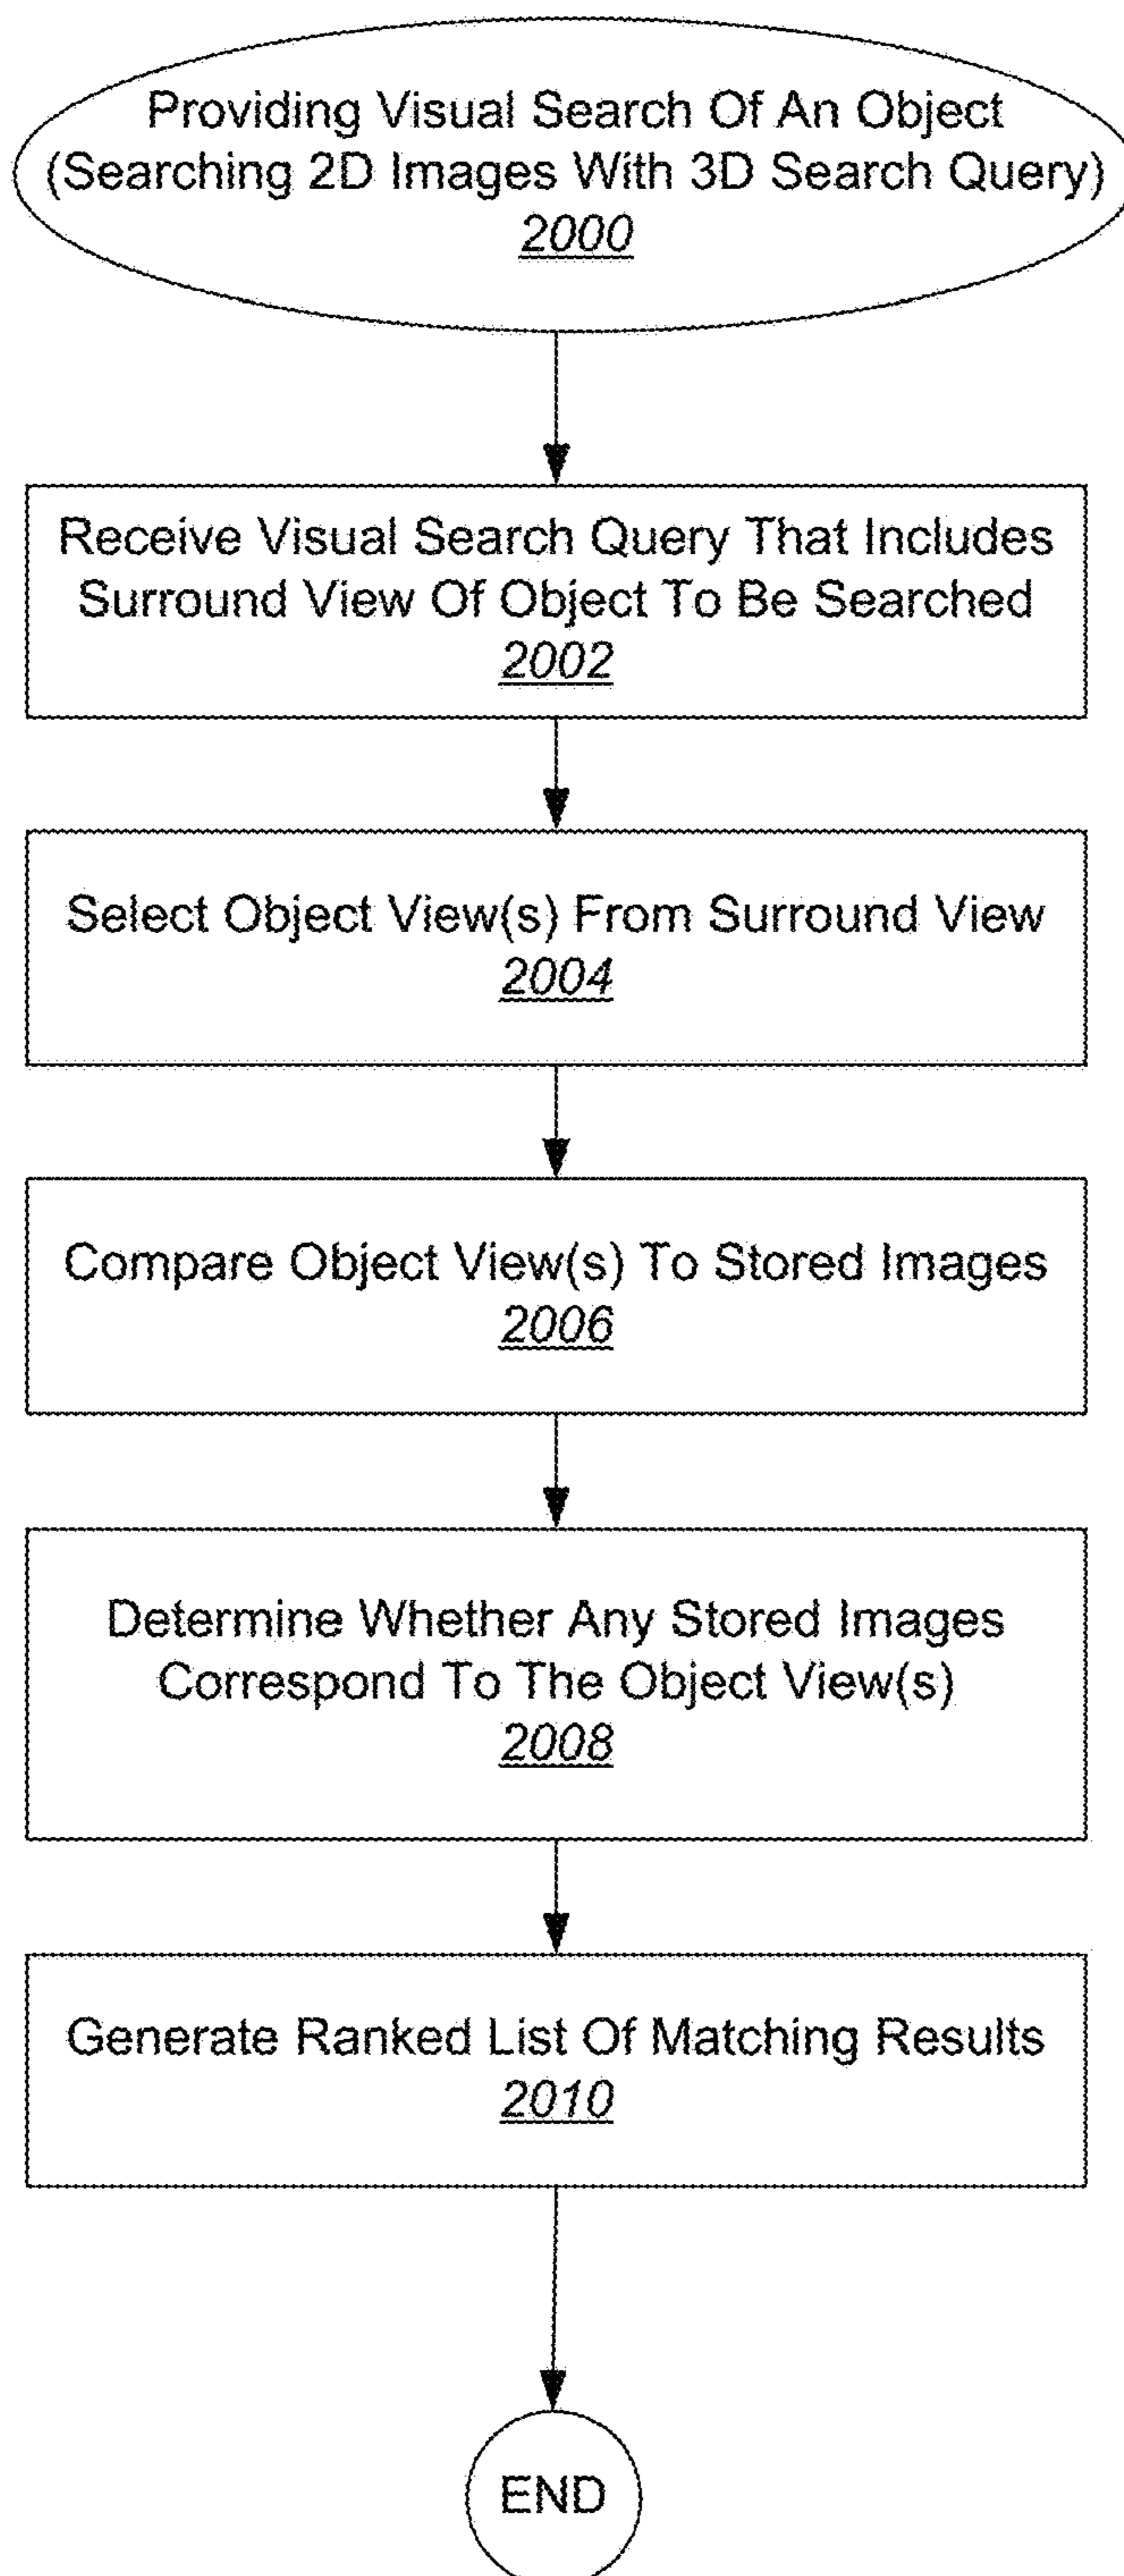


FIG. 20




FIG. 21

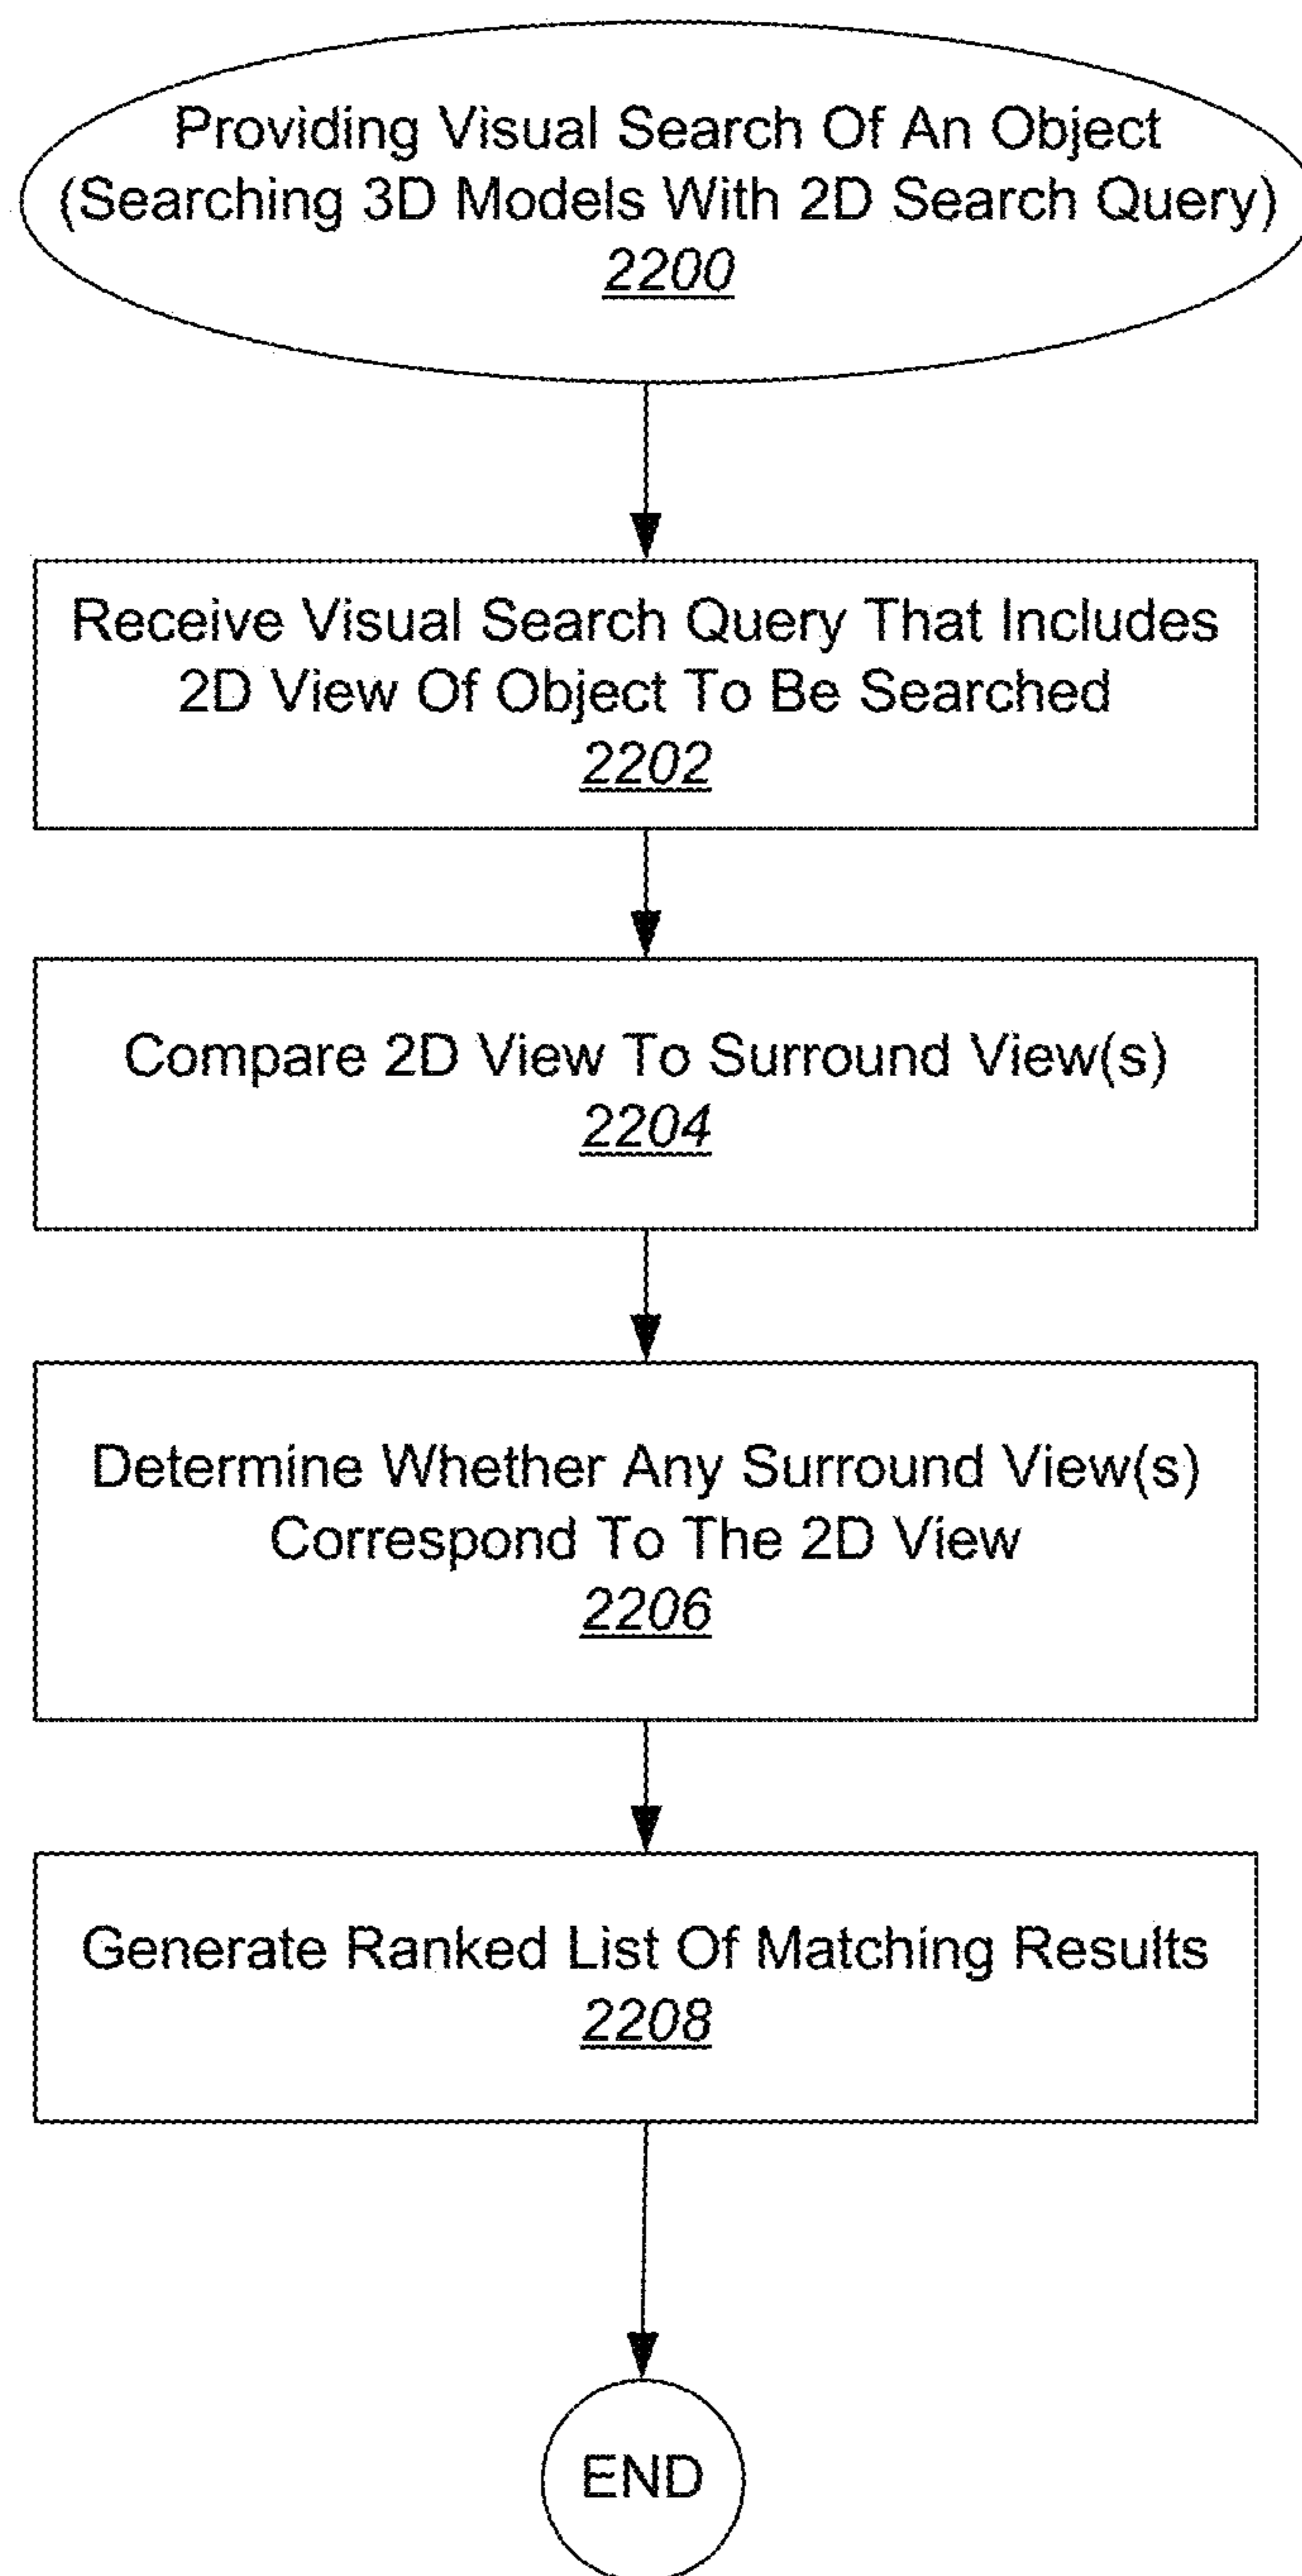


FIG. 22

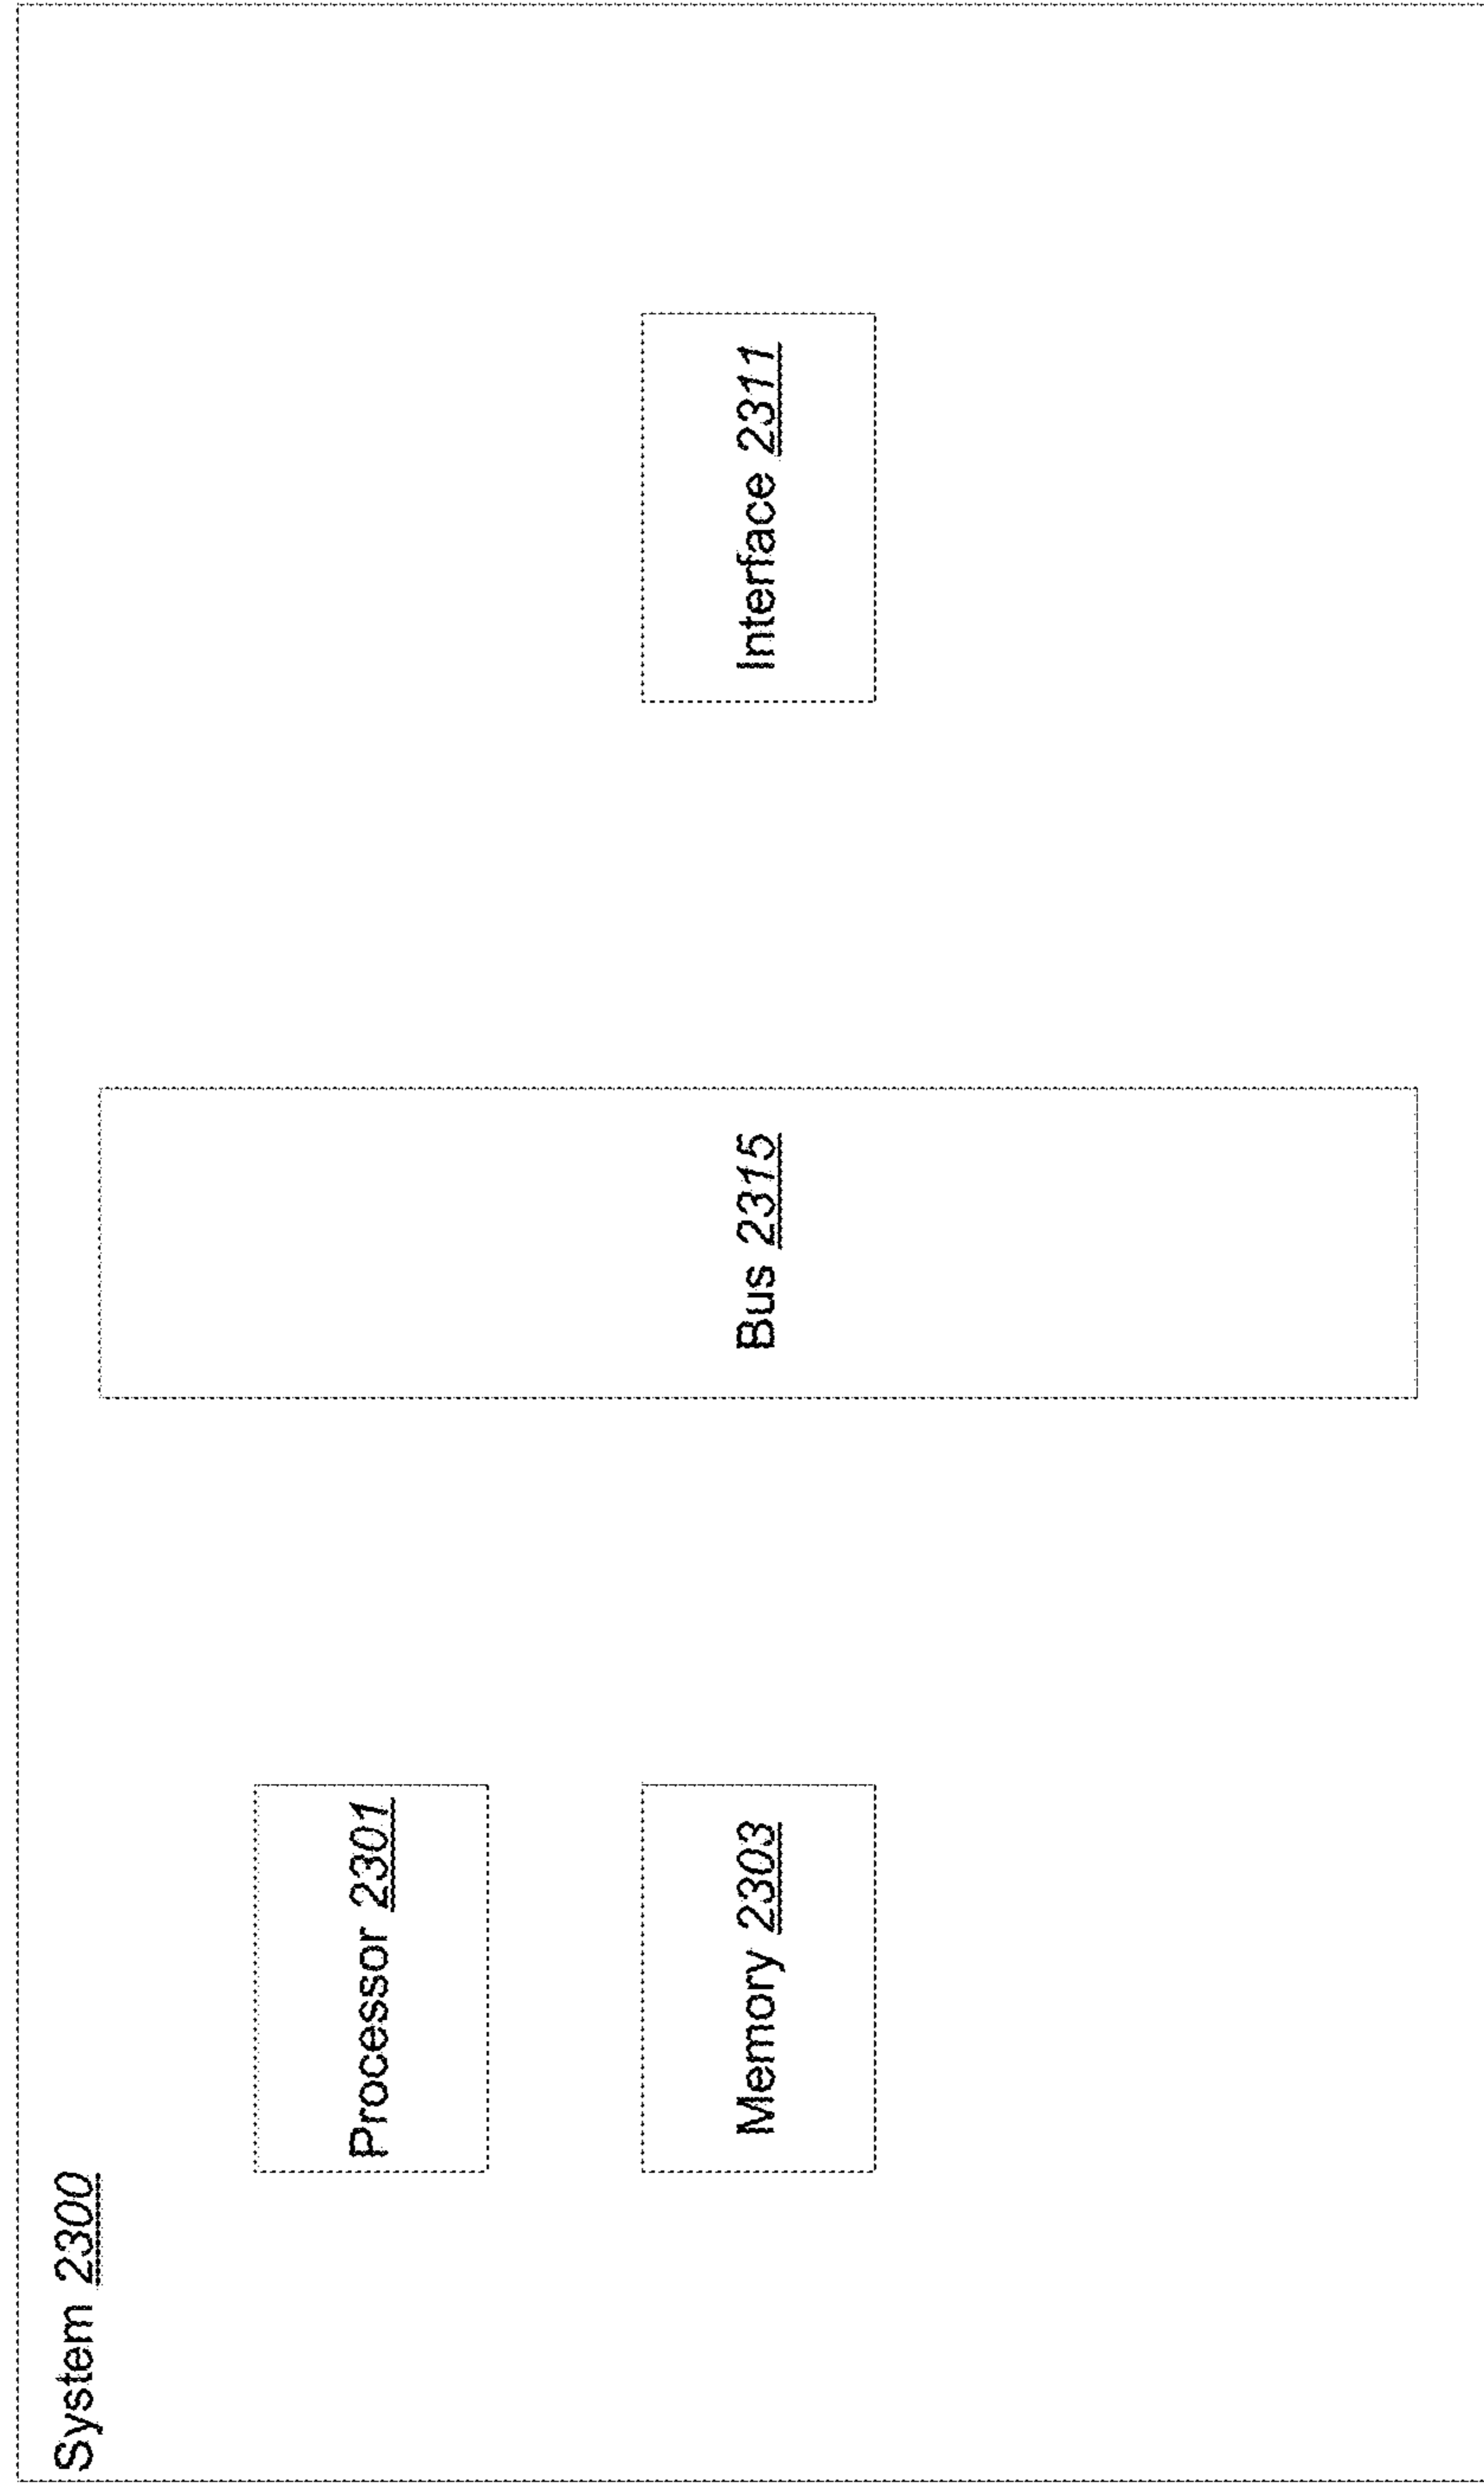


FIG. 23

ANALYSIS AND MANIPULATION OF PANORAMIC SURROUND VIEWS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 14/539,889 by Holzer et al., filed on Nov. 12, 2014, titled "Analysis and Manipulation of Panoramic Surround Views," which claims priority under 35 U.S.C. § 119 to Provisional U.S. Patent Application No. 61/903,359 by Holzer et al., filed on Nov. 12, 2013, titled "Systems and Methods For Providing Surround Views," which are incorporated by reference herein in their entirety and for all purposes.

TECHNICAL FIELD

The present disclosure relates to panoramic surround views, which includes providing a multi-view interactive digital media representation.

DESCRIPTION OF RELATED ART

With modern computing platforms and technologies shifting towards mobile and wearable devices that include camera sensors as native acquisition input streams, the desire to record and preserve moments digitally in a different form than more traditional two-dimensional (2D) flat images and videos has become more apparent. Traditional digital media formats typically limit their viewers to a passive experience. For instance, a 2D flat image can be viewed from one angle and is limited to zooming in and out. Accordingly, traditional digital media formats, such as 2D flat images, do not easily lend themselves to reproducing memories and events with high fidelity.

Current predictions (Ref: KPCB "Internet Trends 2012" presentation") indicate that every several years the quantity of visual data that is being captured digitally online will double. As this quantity of visual data increases, so does the need for much more comprehensive search and indexing mechanisms than ones currently available. Unfortunately, neither 2D images nor 2D videos have been designed for these purposes. Accordingly, improved mechanisms that allow users to view and index visual data, as well as query and quickly receive meaningful results from visual data are desirable.

OVERVIEW

Various embodiments of the present invention relate generally to systems and methods for analyzing and manipulating images and video. According to particular embodiments, the spatial relationship between multiple images and video is analyzed together with location information data, for purposes of creating a representation referred to herein as a surround view. In particular embodiments, a surround view can be generated by combining a panoramic view of an object with a panoramic view of a distant scene, such that the object panorama is placed in a foreground position relative to the distant scene panorama. Such combined panoramas can enhance the interactive and immersive viewing experience of the surround view.

BRIEF DESCRIPTION OF THE DRAWINGS

The disclosure may best be understood by reference to the following description taken in conjunction with the accompanying drawings, which illustrate particular embodiments of the present invention.

FIG. 1 illustrates an example of a surround view acquisition system.

FIG. 2 illustrates an example of a process flow for generating a surround view.

5 FIG. 3 illustrates one example of multiple camera views that can be fused into a three-dimensional (3D) model to create an immersive experience.

FIG. 4A illustrates one example of separation of content and context in a surround view.

10 FIG. 4B illustrates one example of layering in a surround view.

FIG. 4C illustrates one example of a process for modifying a layer in a surround view.

15 FIGS. 5A-5B illustrate examples of concave view and convex views, respectively, where both views use a back-camera capture style.

FIGS. 6A-6E illustrate examples of various capture modes for surround views.

20 FIG. 7A illustrates one example of a process for recording data that can be used to generate a surround view.

FIG. 7B illustrates one example of a dynamic panorama capture process.

25 FIG. 7C illustrates one example of a dynamic panorama capture process where the capture device is rotated through the axis of rotation.

FIG. 7D illustrates one example of a dynamic panorama with dynamic content.

30 FIG. 7E illustrates one example of capturing a dynamic panorama with a 3D effect.

FIG. 7F illustrates one example of a dynamic panorama with parallax effect.

FIG. 7G illustrates one example of an object panorama capture process.

35 FIG. 7H illustrates one example of a background panorama with an object panorama projected on it.

FIG. 7I illustrates one example of multiple objects constituting an object panorama.

40 FIG. 7J illustrates one example of changing the viewing angle of an object panorama based on user navigation.

FIG. 7K illustrates one example of a selfie panorama capture process.

FIG. 7L illustrates one example of a background panorama with a selfie panorama projected on it.

45 FIG. 7M illustrates one example of extended views of panoramas based on user navigation.

FIG. 8 illustrates an example of a surround view in which three-dimensional content is blended with a two-dimensional panoramic context.

50 FIG. 9 illustrates one example of a space-time surround view being simultaneously recorded by independent observers.

FIG. 10 illustrates one example of separation of a complex surround-view into smaller, linear parts.

55 FIG. 11 illustrates one example of a combination of multiple surround views into a multi-surround view.

FIG. 12 illustrates one example of a process for prompting a user for additional views of an object of interest to provide a more accurate surround view.

60 FIGS. 13A-13B illustrate an example of prompting a user for additional views of an object to be searched.

FIG. 14 illustrates one example of a process for navigating a surround view.

65 FIG. 15 illustrates an example of swipe-based navigation of a surround view.

FIG. 16A illustrates examples of a sharing service for surround views, as shown on a mobile device and browser.

FIG. 16B illustrates examples of surround view-related notifications on a mobile device.

FIG. 17A illustrates one example of a process for providing object segmentation.

FIG. 17B illustrates one example of a segmented object viewed from different angles.

FIG. 18 illustrates one example of various data sources that can be used for surround view generation and various applications that can be used with a surround view.

FIG. 19 illustrates one example of a process for providing visual search of an object, where the search query includes a surround view of the object and the data searched includes three-dimensional models.

FIG. 20 illustrates one example of a process for providing visual search of an object, where the search query includes a surround view of the object and the data searched includes two-dimensional images.

FIG. 21 illustrates an example of a visual search process.

FIG. 22 illustrates an example of a process for providing visual search of an object, where the search query includes a two-dimensional view of the object and the data searched includes surround view(s).

FIG. 23 illustrates a particular example of a computer system that can be used with various embodiments of the present invention.

DETAILED DESCRIPTION

Reference will now be made in detail to some specific examples of the invention including the best modes contemplated by the inventors for carrying out the invention. Examples of these specific embodiments are illustrated in the accompanying drawings. While the present disclosure is described in conjunction with these specific embodiments, it will be understood that it is not intended to limit the invention to the described embodiments. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims.

In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. Particular embodiments of the present invention may be implemented without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

Various aspects of the present invention relate generally to systems and methods for analyzing the spatial relationship between multiple images and video together with location information data, for the purpose of creating a single representation, a surround view, which eliminates redundancy in the data, and presents a user with an interactive and immersive active viewing experience. According to various embodiments, active is described in the context of providing a user with the ability to control the viewpoint of the visual information displayed on a screen. In particular example embodiments, the surround view data structure (and associated algorithms) is natively built for, but not limited to, applications involving visual search.

According to various embodiments of the present invention, a surround view is a multi-view interactive digital media representation. With reference to FIG. 1, shown is one example of a surround view acquisition system 100. In the present example embodiment, the surround view acquisition system 100 is depicted in a flow sequence that can be used to generate a surround view. According to various embodiments, the data used to generate a surround view can come

from a variety of sources. In particular, data such as, but not limited to two-dimensional (2D) images 104 can be used to generate a surround view. These 2D images can include color image data streams such as multiple image sequences, video data, etc., or multiple images in any of various formats for images, depending on the application. Another source of data that can be used to generate a surround view includes location information 106. This location information 106 can be obtained from sources such as accelerometers, gyroscopes, magnetometers, GPS, WiFi, IMU-like systems (Inertial Measurement Unit systems), and the like. Yet another source of data that can be used to generate a surround view can include depth images 108. These depth images can include depth, 3D, or disparity image data streams, and the like, and can be captured by devices such as, but not limited to, stereo cameras, time-of-flight cameras, three-dimensional cameras, and the like.

In the present example embodiment, the data can then be fused together at sensor fusion block 110. In some embodiments, a surround view can be generated a combination of data that includes both 2D images 104 and location information 106, without any depth images 108 provided. In other embodiments, depth images 108 and location information 106 can be used together at sensor fusion block 110. Various combinations of image data can be used with location information at 106, depending on the application and available data.

In the present example embodiment, the data that has been fused together at sensor fusion block 110 is then used for content modeling 112 and context modeling 114. As described in more detail with regard to FIG. 4, the subject matter featured in the images can be separated into content and context. The content can be delineated as the object of interest and the context can be delineated as the scenery surrounding the object of interest. According to various embodiments, the content can be a three-dimensional model, depicting an object of interest, although the content can be a two-dimensional image in some embodiments, as described in more detail below with regard to FIG. 4. Furthermore, in some embodiments, the context can be a two-dimensional model depicting the scenery surrounding the object of interest. Although in many examples the context can provide two-dimensional views of the scenery surrounding the object of interest, the context can also include three-dimensional aspects in some embodiments. For instance, the context can be depicted as a “flat” image along a cylindrical “canvas,” such that the “flat” image appears on the surface of a cylinder. In addition, some examples may include three-dimensional context models, such as when some objects are identified in the surrounding scenery as three-dimensional objects. According to various embodiments, the models provided by content modeling 112 and context modeling 114 can be generated by combining the image and location information data, as described in more detail with regard to FIG. 3.

According to various embodiments, context and content of a surround view are determined based on a specified object of interest. In some examples, an object of interest is automatically chosen based on processing of the image and location information data. For instance, if a dominant object is detected in a series of images, this object can be selected as the content. In other examples, a user specified target 102 can be chosen, as shown in FIG. 1. It should be noted, however, that a surround view can be generated without a user specified target in some applications.

In the present example embodiment, one or more enhancement algorithms can be applied at enhancement

algorithm(s) block **116**. In particular example embodiments, various algorithms can be employed during capture of surround view data, regardless of the type of capture mode employed. These algorithms can be used to enhance the user experience. For instance, automatic frame selection, stabilization, view interpolation, filters, and/or compression can be used during capture of surround view data. In some examples, these enhancement algorithms can be applied to image data after acquisition of the data. In other examples, these enhancement algorithms can be applied to image data during capture of surround view data.

According to particular example embodiments, automatic frame selection can be used to create a more enjoyable surround view. Specifically, frames are automatically selected so that the transition between them will be smoother or more even. This automatic frame selection can incorporate blur- and overexposure-detection in some applications, as well as more uniformly sampling poses such that they are more evenly distributed.

In some example embodiments, stabilization can be used for a surround view in a manner similar to that used for video. In particular, keyframes in a surround view can be stabilized for to produce improvements such as smoother transitions, improved/enhanced focus on the content, etc. However, unlike video, there are many additional sources of stabilization for a surround view, such as by using IMU information, depth information, computer vision techniques, direct selection of an area to be stabilized, face detection, and the like.

For instance, IMU information can be very helpful for stabilization. In particular, IMU information provides an estimate, although sometimes a rough or noisy estimate, of the camera tremor that may occur during image capture. This estimate can be used to remove, cancel, and/or reduce the effects of such camera tremor.

In some examples, depth information, if available, can be used to provide stabilization for a surround view. Because points of interest in a surround view are three-dimensional, rather than two-dimensional, these points of interest are more constrained and tracking/matching of these points is simplified as the search space reduces. Furthermore, descriptors for points of interest can use both color and depth information and therefore, become more discriminative. In addition, automatic or semi-automatic content selection can be easier to provide with depth information. For instance, when a user selects a particular pixel of an image, this selection can be expanded to fill the entire surface that touches it. Furthermore, content can also be selected automatically by using a foreground/background differentiation based on depth. In various examples, the content can stay relatively stable/visible even when the context changes.

According to various examples, computer vision techniques can also be used to provide stabilization for surround views. For instance, keypoints can be detected and tracked. However, in certain scenes, such as a dynamic scene or static scene with parallax, no simple warp exists that can stabilize everything. Consequently, there is a trade-off in which certain aspects of the scene receive more attention to stabilization and other aspects of the scene receive less attention. Because a surround view is often focused on a particular object of interest, a surround view can be content-weighted so that the object of interest is maximally stabilized in some examples.

Another way to improve stabilization in a surround view includes direct selection of a region of a screen. For instance, if a user taps to focus on a region of a screen, then records a convex surround view, the area that was tapped can be

maximally stabilized. This allows stabilization algorithms to be focused on a particular area or object of interest.

In some examples, face detection can be used to provide stabilization. For instance, when recording with a front-facing camera, it is often likely that the user is the object of interest in the scene. Thus, face detection can be used to weight stabilization about that region. When face detection is precise enough, facial features themselves (such as eyes, nose, mouth) can be used as areas to stabilize, rather than using generic keypoints.

According to various examples, view interpolation can be used to improve the viewing experience. In particular, to avoid sudden “jumps” between stabilized frames, synthetic, intermediate views can be rendered on the fly. This can be informed by content-weighted keypoint tracks and IMU information as described above, as well as by denser pixel-to-pixel matches. If depth information is available, fewer artifacts resulting from mismatched pixels may occur, thereby simplifying the process. As described above, view interpolation can be applied during capture of a surround view in some embodiments. In other embodiments, view interpolation can be applied during surround view generation.

In some examples, filters can also be used during capture or generation of a surround view to enhance the viewing experience. Just as many popular photo sharing services provide aesthetic filters that can be applied to static, two-dimensional images, aesthetic filters can similarly be applied to surround images. However, because a surround view representation is more expressive than a two-dimensional image, and three-dimensional information is available in a surround view, these filters can be extended to include effects that are ill-defined in two dimensional photos. For instance, in a surround view, motion blur can be added to the background (i.e. context) while the content remains crisp. In another example, a drop-shadow can be added to the object of interest in a surround view.

In various examples, compression can also be used as an enhancement algorithm **116**. In particular, compression can be used to enhance user-experience by reducing data upload and download costs. Because surround views use spatial information, far less data can be sent for a surround view than a typical video, while maintaining desired qualities of the surround view. Specifically, the IMU, keypoint tracks, and user input, combined with the view interpolation described above, can all reduce the amount of data that must be transferred to and from a device during upload or download of a surround view. For instance, if an object of interest can be properly identified, a variable compression style can be chosen for the content and context. This variable compression style can include lower quality resolution for background information (i.e. context) and higher quality resolution for foreground information (i.e. content) in some examples. In such examples, the amount of data transmitted can be reduced by sacrificing some of the context quality, while maintaining a desired level of quality for the content.

In the present embodiment, a surround view **118** is generated after any enhancement algorithms are applied. The surround view can provide a multi-view interactive digital media representation. In various examples, the surround view can include three-dimensional model of the content and a two-dimensional model of the context. However, in some examples, the context can represent a “flat” view of the scenery or background as projected along a surface, such as a cylindrical or other-shaped surface, such

that the context is not purely two-dimensional. In yet other examples, the context can include three-dimensional aspects.

According to various embodiments, surround views provide numerous advantages over traditional two-dimensional images or videos. Some of these advantages include: the ability to cope with moving scenery, a moving acquisition device, or both; the ability to model parts of the scene in three-dimensions; the ability to remove unnecessary, redundant information and reduce the memory footprint of the output dataset; the ability to distinguish between content and context; the ability to use the distinction between content and context for improvements in the user-experience; the ability to use the distinction between content and context for improvements in memory footprint (an example would be high quality compression of content and low quality compression of context); the ability to associate special feature descriptors with surround views that allow the surround views to be indexed with a high degree of efficiency and accuracy; and the ability of the user to interact and change the viewpoint of the surround view. In particular example embodiments, the characteristics described above can be incorporated natively in the surround view representation, and provide the capability for use in various applications. For instance, surround views can be used to enhance various fields such as e-commerce, visual search, 3D printing, file sharing, user interaction, and entertainment.

According to various example embodiments, once a surround view **118** is generated, user feedback for acquisition **120** of additional image data can be provided. In particular, if a surround view is determined to need additional views to provide a more accurate model of the content or context, a user may be prompted to provide additional views. Once these additional views are received by the surround view acquisition system **100**, these additional views can be processed by the system **100** and incorporated into the surround view.

With reference to FIG. 2, shown is an example of a process flow diagram for generating a surround view **200**. In the present example, a plurality of images is obtained at **202**. According to various embodiments, the plurality of images can include two-dimensional (2D) images or data streams. These 2D images can include location information that can be used to generate a surround view. In some embodiments, the plurality of images can include depth images **108**, as also described above with regard to FIG. 1. The depth images can also include location information in various examples.

According to various embodiments, the plurality of images obtained at **202** can include a variety of sources and characteristics. For instance, the plurality of images can be obtained from a plurality of users. These images can be a collection of images gathered from the internet from different users of the same event, such as 2D images or video obtained at a concert, etc. In some examples, the plurality of images can include images with different temporal information. In particular, the images can be taken at different times of the same object of interest. For instance, multiple images of a particular statue can be obtained at different times of day, different seasons, etc. In other examples, the plurality of images can represent moving objects. For instance, the images may include an object of interest moving through scenery, such as a vehicle traveling along a road or a plane traveling through the sky. In other instances, the images may include an object of interest that is also moving, such as a person dancing, running, twirling, etc.

In the present example embodiment, the plurality of images is fused into content and context models at **204**.

According to various embodiments, the subject matter featured in the images can be separated into content and context. The content can be delineated as the object of interest and the context can be delineated as the scenery surrounding the object of interest. According to various embodiments, the content can be a three-dimensional model, depicting an object of interest, and the content can be a two-dimensional image in some embodiments.

According to the present example embodiment, one or more enhancement algorithms can be applied to the content and context models at **206**. These algorithms can be used to enhance the user experience. For instance, enhancement algorithms such as automatic frame selection, stabilization, view interpolation, filters, and/or compression can be used. In some examples, these enhancement algorithms can be applied to image data during capture of the images. In other examples, these enhancement algorithms can be applied to image data after acquisition of the data.

In the present embodiment, a surround view is generated from the content and context models at **208**. The surround view can provide a multi-view interactive digital media representation. In various examples, the surround view can include a three-dimensional model of the content and a two-dimensional model of the context. According to various embodiments, depending on the mode of capture and the viewpoints of the images, the surround view model can include certain characteristics. For instance, some examples of different styles of surround views include a locally concave surround view, a locally convex surround view, and a locally flat surround view. However, it should be noted that surround views can include combinations of views and characteristics, depending on the application.

With reference to FIG. 3, shown is one example of multiple camera views that can be fused together into a three-dimensional (3D) model to create an immersive experience. According to various embodiments, multiple images can be captured from various viewpoints and fused together to provide a surround view. In the present example embodiment, three cameras **312**, **314**, and **316** are positioned at locations **322**, **324**, and **326**, respectively, in proximity to an object of interest **308**. Scenery can surround the object of interest **308** such as object **310**. Views **302**, **304**, and **306** from their respective cameras **312**, **314**, and **316** include overlapping subject matter. Specifically, each view **302**, **304**, and **306** includes the object of interest **308** and varying degrees of visibility of the scenery surrounding the object **310**. For instance, view **302** includes a view of the object of interest **308** in front of the cylinder that is part of the scenery surrounding the object **310**. View **306** shows the object of interest **308** to one side of the cylinder, and view **304** shows the object of interest without any view of the cylinder.

In the present example embodiment, the various views **302**, **304**, and **316** along with their associated locations **322**, **324**, and **326**, respectively, provide a rich source of information about object of interest **308** and the surrounding context that can be used to produce a surround view. For instance, when analyzed together, the various views **302**, **304**, and **326** provide information about different sides of the object of interest and the relationship between the object of interest and the scenery. According to various embodiments, this information can be used to parse out the object of interest **308** into content and the scenery as the context. Furthermore, as also described above with regard to FIGS. **1** and **2**, various algorithms can be applied to images produced by these viewpoints to create an immersive, interactive experience when viewing a surround view.

FIG. 4A illustrates one example of separation of content and context in a surround view. According to various embodiments of the present invention, a surround view is a multi-view interactive digital media representation of a scene 400. With reference to FIG. 4A, shown is a user 402 located in a scene 400. The user 402 is capturing images of an object of interest, such as a statue. The images captured by the user constitute digital visual data that can be used to generate a surround view.

According to various embodiments of the present disclosure, the digital visual data included in a surround view can be, semantically and/or practically, separated into content 404 and context 406. According to particular embodiments, content 404 can include the object(s), person(s), or scene(s) of interest while the context 406 represents the remaining elements of the scene surrounding the content 404. In some examples, a surround view may represent the content 404 as three-dimensional data, and the context 406 as a two-dimensional panoramic background. In other examples, a surround view may represent both the content 404 and context 406 as two-dimensional panoramic scenes. In yet other examples, content 404 and context 406 may include three-dimensional components or aspects. In particular embodiments, the way that the surround view depicts content 404 and context 406 depends on the capture mode used to acquire the images.

In some examples, such as but not limited to: recordings of objects, persons, or parts of objects or persons, where only the object, person, or parts of them are visible, recordings of large flat areas, and recordings of scenes where the data captured appears to be at infinity (i.e., there are no subjects close to the camera), the content 404 and the context 406 may be the same. In these examples, the surround view produced may have some characteristics that are similar to other types of digital media such as panoramas. However, according to various embodiments, surround views include additional features that distinguish them from these existing types of digital media. For instance, a surround view can represent moving data. Additionally, a surround view is not limited to a specific cylindrical, spherical or translational movement. Various motions can be used to capture image data with a camera or other capture device. Furthermore, unlike a stitched panorama, a surround view can display different sides of the same object.

Although a surround view can be separated into content and context in some applications, a surround view can also be separated into layers in other applications. With reference to FIG. 4B, shown is one example of layering in a surround view. In this example, a layered surround view 410 is segmented into different layers 418, 420, and 422. Each layer 418, 420, and 422 can include an object (or a set of objects), people, dynamic scene elements, background, etc. Furthermore, each of these layers 418, 420, and 422 can be assigned a depth.

According to various embodiments, the different layers 418, 420, and 422 can be displayed in different ways. For instance, different filters (e.g. gray scale filter, blurring, etc.) can be applied to some layers but not to others. In other examples, different layers can be moved at different speeds relative to each other, such that when a user swipes through a surround view a better three-dimensional effect is provided. Similarly, when a user swipes along the parallax direction, the layers can be displaced differently to provide a better three-dimensional effect. In addition, one or more layers can be omitted when displaying a surround view, such that unwanted objects, etc. can be removed from a surround view.

In the present example, a user 412 is shown holding a capture device 414. The user 412 moves the capture device 414 along capture motion 416. When the images captured are used to generate a surround view, layers 418, 420, and 422 are separated based on depth. These layers can then be processed or viewed differently in a surround view, depending on the application.

With reference to FIG. 4C, shown is one example of a process for generating a surround view with a modified layer in a surround view 430. In particular, a first surround view having a first layer and a second layer is obtained at 432. As described above with regard to FIG. 4B, a surround view can be divided into different layers. In the present example, the first layer includes a first depth and the second layer includes a second depth.

Next, the first layer is selected at 434. According to various examples, selecting the first layer includes selecting data within the first depth. More specifically, selecting data within the first depth includes selecting the visual data located within the first depth. According to various embodiments, the first layer can include features such as an object, person, dynamic scene elements, background, etc. In some examples, selection of the first layer is performed automatically without user input. In other examples, selection of the first layer is performed semi-automatically using user-guided interaction.

After the first layer is selected, an effect is applied to the first layer within the first surround view to produce a modified first layer at 436. In one example, the effect applied can be a filter such as a blurring filter, gray scale filter, etc. In another example, the effect applied can include moving the first layer at a first speed relative to the second layer, which is moved at a second speed. When the first speed is different from the second speed, three-dimensional effects can be improved in some instances. In some applications, a parallax effect can occur, thereby creating a three-dimensional effect.

Next, a second surround view is generated that includes the modified first layer and the second layer at 438. As described above, applying one or more effects to the first layer can improve the three-dimensional effects of a surround view in some applications. In these applications, the second surround view can have improved three-dimensional effects when compared to the first surround view. Other effects can be applied in different examples, and can emphasize or deemphasize various aspects of a first surround view to yield a second surround view. In addition, in some applications, a layer can be omitted in a second surround view. Specifically, when the first surround view includes a third layer, the second surround view omits this third layer. In one example, this third layer could include an object or person that would be "edited out" in the generated second surround view. In another example, this third layer could include a background or background elements, and the second surround view generated would not include the background or background elements. Of course, any object or feature can be located in this omitted third layer, depending on the application.

FIGS. 5A-5B illustrate examples of concave and convex views, respectively, where both views use a back-camera capture style. In particular, if a camera phone is used, these views use the camera on the back of the phone, facing away from the user. In particular embodiments, concave and convex views can affect how the content and context are designated in a surround view.

With reference to FIG. 5A, shown is one example of a concave view 500 in which a user is standing along a vertical

axis **508**. In this example, the user is holding a camera, such that camera location **502** does not leave axis **508** during image capture. However, as the user pivots about axis **508**, the camera captures a panoramic view of the scene around the user, forming a concave view. In this embodiment, the object of interest **504** and the distant scenery **506** are all viewed similarly because of the way in which the images are captured. In this example, all objects in the concave view appear at infinity, so the content is equal to the context according to this view.

With reference to FIG. **5B**, shown is one example of a convex view **520** in which a user changes position when capturing images of an object of interest **524**. In this example, the user moves around the object of interest **524**, taking pictures from different sides of the object of interest from camera locations **528**, **530**, and **532**. Each of the images obtained includes a view of the object of interest, and a background of the distant scenery **526**. In the present example, the object of interest **524** represents the content, and the distant scenery **526** represents the context in this convex view.

FIGS. **6A-6E** illustrate examples of various capture modes for surround views. Although various motions can be used to capture a surround view and are not constrained to any particular type of motion, three general types of motion can be used to capture particular features or views described in conjunction surround views. These three types of motion, respectively, can yield a locally concave surround view, a locally convex surround view, and a locally flat surround view. In some examples, a surround view can include various types of motions within the same surround view.

With reference to FIG. **6A**, shown is an example of a back-facing, concave surround view being captured. According to various embodiments, a locally concave surround view is one in which the viewing angles of the camera or other capture device diverge. In one dimension this can be likened to the motion required to capture a spherical 360 panorama (pure rotation), although the motion can be generalized to any curved sweeping motion in which the view faces outward. In the present example, the experience is that of a stationary viewer looking out at a (possibly dynamic) context.

In the present example embodiment, a user **602** is using a back-facing camera **606** to capture images towards world **600**, and away from user **602**. As described in various examples, a back-facing camera refers to a device with a camera that faces away from the user, such as the camera on the back of a smart phone. The camera is moved in a concave motion **608**, such that views **604a**, **604b**, and **604c** capture various parts of capture area **609**.

With reference to FIG. **6B**, shown is an example of a back-facing, convex surround view being captured. According to various embodiments, a locally convex surround view is one in which viewing angles converge toward a single object of interest. In some examples, a locally convex surround view can provide the experience of orbiting about a point, such that a viewer can see multiple sides of the same object. This object, which may be an “object of interest,” can be segmented from the surround view to become the content, and any surrounding data can be segmented to become the context. Previous technologies fail to recognize this type of viewing angle in the media-sharing landscape.

In the present example embodiment, a user **602** is using a back-facing camera **614** to capture images towards world **600**, and away from user **602**. The camera is moved in a convex motion **610**, such that views **612a**, **612b**, and **612c** capture various parts of capture area **611**. As described

above, world **600** can include an object of interest in some examples, and the convex motion **610** can orbit around this object. Views **612a**, **612b**, and **612c** can include views of different sides of this object in these examples.

With reference to FIG. **6C**, shown is an example of a front-facing, concave surround view being captured. As described in various examples, a front-facing camera refers to a device with a camera that faces towards the user, such as the camera on the front of a smart phone. For instance, front-facing cameras are commonly used to take “selfies” (i.e., self-portraits of the user).

In the present example embodiment, camera **620** is facing user **602**. The camera follows a concave motion **606** such that the views **618a**, **618b**, and **618c** diverge from each other in an angular sense. The capture area **617** follows a concave shape that includes the user at a perimeter.

With reference to FIG. **6D**, shown is an example of a front-facing, convex surround view being captured. In the present example embodiment, camera **626** is facing user **602**. The camera follows a convex motion **622** such that the views **624a**, **624b**, and **624c** converge towards the user **602**. The capture area **617** follows a concave shape that surrounds the user **602**.

With reference to FIG. **6E**, shown is an example of a back-facing, flat view being captured. In particular example embodiments, a locally flat surround view is one in which the rotation of the camera is small compared to its translation. In a locally flat surround view, the viewing angles remain roughly parallel, and the parallax effect dominates. In this type of surround view, there can also be an “object of interest”, but its position does not remain fixed in the different views. Previous technologies also fail to recognize this type of viewing angle in the media-sharing landscape.

In the present example embodiment, camera **632** is facing away from user **602**, and towards world **600**. The camera follows a generally linear motion **628** such that the capture area **629** generally follows a line. The views **630a**, **630b**, and **630c** have generally parallel lines of sight. An object viewed in multiple views can appear to have different or shifted background scenery in each view. In addition, a slightly different side of the object may be visible in different views. Using the parallax effect, information about the position and characteristics of the object can be generated in a surround view that provides more information than any one static image.

As described above, various modes can be used to capture images for a surround view. These modes, including locally concave, locally convex, and locally linear motions, can be used during capture of separate images or during continuous recording of a scene. Such recording can capture a series of images during a single session.

According to various embodiments of the present invention, a surround view can be generated from data acquired in numerous ways. FIG. **7A** illustrates one example of process for recording data that can be used to generate a surround view. In this example, data is acquired by moving a camera through space. In particular, a user taps a record button **702** on a capture device **700** to begin recording. As movement of the capture device **716** follows a generally leftward direction, an object **714** moves in a generally rightward motion across the screen, as indicated by movement of object **716**. Specifically, the user presses the record button **702** in view **708**, and then moves the capture device leftward in view **710**. As the capture device moves leftward, object **714** appears to move rightward between views **710** and **712**. In some examples, when the user is finished recording, the record button **702** can be tapped again. In

other examples, the user can tap and hold the record button during recording, and release to stop recording. In the present embodiment, the recording captures a series of images that can be used to generate a surround view.

According to various embodiments, different types of panoramas can be captured in surround views, depending on the type of movement used in the capture process. In particular, dynamic panoramas, object panoramas, and selfie panoramas can be generated based on captured data. In some embodiments, the captured data can be recorded as described with regard to FIG. 7A.

FIGS. 7B-7F illustrate examples relating to dynamic panoramas that can be created with surround views. With particular reference to FIG. 7B, shown is one example of a dynamic panorama capture process 720. In the present example, a user 722 moves capture device 724 along capture motion 726. This capture motion 726 can include rotating, waving, translating, etc. the capture device 724. During this capture process, a panorama of scene 728 is generated and dynamic content within the scene is kept. For instance, moving objects are preserved within the panorama as dynamic content.

With reference to FIG. 7C, shown is a specific example of a dynamic panorama capture process 730 where a capture device 732 is rotated through an axis of rotation 734. In particular, capture device 732 is rotated about its center along an axis of rotation 734. This pure rotation captures a panorama of scene 736. According to various examples, this type of panorama can provide a “flat” scene that captures entities in the scene at a particular point in time. This “flat” scene can be a two-dimensional image, or can be an image projected on a cylinder, surface, etc.

With reference to FIG. 7D, shown is one example of a dynamic panorama 740 with dynamic content 744. Once a panorama is captured, as described above with regard to FIGS. 7B-7C, a dynamic panorama 740 can be navigated by a user. In the present example, dynamic content 744 is animated when the user navigates through the dynamic panorama 740. For instance, as the user swipes across scene 742, the dynamic content 744 can be seen moving with respect to the scene 742.

With reference to FIG. 7E, shown is one example of capturing a dynamic panorama with a 3D effect. In the present example, if a capture device is not rotated exactly around its camera center (as in FIG. 7C), a 3D effect can be obtained by moving different parts of the panorama at different speeds while the user navigates through the dynamic content. Although a nearby person or object 750 would create artifacts in a standard panorama capture process if the capture device is not rotated around its camera center (as in FIG. 7C), these “imperfections” can be used to create a 3D impression to the user by moving the object 750 at a different speed when swiping/navigating through a dynamic panorama. In particular, the capture device 745 shown uses a capture motion 748 that captures a distant scene 746 and a nearby person/object 750. The movements of the nearby person/object 750 can be captured as 3D motion within the surround view, while the distant scenery 746 appears to be static as the user navigates through the surround view, according to various embodiments.

With reference to FIG. 7F, shown is one example of a dynamic panorama 750 with parallax effect. Three-dimensional effects can be presented by applying a parallax effect when swiping perpendicular to the panorama direction 752. In particular, when swiping perpendicular to the panorama direction, along the parallax direction 754, nearby objects

are displaced along the parallax direction 754 while the scene at distance stays still or moves less than the nearby objects.

FIGS. 7G-7J illustrate examples relating to object panoramas that can be created with surround views. With reference to FIG. 7G, shown is one example of an object panorama capture process. In particular, a capture device 766 is moved around an object 762 along a capture motion 760. One particular example of a capture device 766 is a smartphone. The capture device 766 also captures a panoramic view of the background 764 as various views and angles of the object 762 are captured. The resulting surround view includes a panoramic view of object 762.

In some embodiments, a surround view can be created by projecting an object panorama onto a background panorama, an example of which is shown in FIG. 7H. In particular, a panorama 768 of this kind is built using background panorama 770 and projecting a foreground object panorama 772 onto the background panorama 770. In some examples, an object panorama can be segmented content taken from a surround view, as described in more detail with regard to FIGS. 17A-17B.

According to various embodiments, multiple objects can make up an object panorama. With reference to FIG. 7I, shown is one example of a capture process for a group of objects 780 making up an object panorama. As shown, a capture device 776 can move around a foreground object, which can be a single object or a group of objects 780 located at a similar distance to the capture device. The capture device 776 can move around the object or group of objects 780 along a capture motion 778, such that various views and angles of the objects are captured. The resulting surround view can include an object panorama of the group of objects 780 with distant background 782 as the context.

Object panoramas allow users to navigate around the object, according to various examples. With reference to FIG. 7J, shown is one example of changing the viewing angle of an object panorama based on user navigation. In this example, three views are shown of a surround view panorama 784. In the surround view panorama, a foreground object 786 is shown in front of a background panorama 788. As a user navigates the panorama by swiping or otherwise interacting with the surround view, the location of the object, the viewing angle of the object, or both can be changed. In the present example, the user can swipe in the direction of the main panorama axis. This navigation can rotate the foreground object 786 in this view. In some examples, the distant background panorama 788 may not change as the foreground object panorama rotates or otherwise moves.

According to various embodiments, object panoramas can also include parallax effects. These parallax effects can be seen when swiping/navigating perpendicular to the direction of the main panorama axis. Similar to FIG. 7F, three-dimensional effects can be presented when swiping perpendicular to the panorama direction. In particular, when swiping perpendicular to the panorama direction, along the parallax direction, nearby objects are displaced along the parallax direction while the scene at distance stays still or moves less than the nearby objects.

Although the previous examples relate to static content and background context in object panoramas, dynamic content can be integrated in the object panorama for either or both the foreground object and the background context. For instance, dynamic content can be featured in a manner similar to that described in conjunction with FIG. 7D. Similarly, dynamic context can also be included in object panoramas.

Another type of panorama that can be included in surround views is a selfie panorama. In some examples, a selfie panorama can be segmented content taken from a surround view, as described in more detail with regard to FIGS. 17A-17B. FIGS. 7K-7L illustrate examples relating to selfie panoramas that can be created with surround views. With reference to FIG. 7K, shown is one example of a selfie panorama capture process 790. In particular, a user 794 moves a capture device 792 along capture motion 796 while capturing images of the user 794. In some examples, the capture device 792 can use a front-facing camera, such as one included on a smart phone. In other examples, a digital camera or other image recording device can be used. A selfie panorama is created with these images, with background 798 providing the context.

With reference to FIG. 7L, shown is one example of a background panorama with a selfie panorama projected on it. In the present example, a surround view panorama 723 is built from a background panorama 725 with a selfie panorama 721 projected on it. According to various examples, the selfie panorama can include a single person or multiple people, similar to the object or group of objects described in conjunction with FIG. 7I. In the present example, selfie panoramas can include dynamic content. For instance, the user can look at the capture device as the capture device moves or the user can keep still while moving the capture device. The user's movements can be captured while the selfie panorama 721 is recorded. These dynamic elements will be mapped into the panorama and can be displayed while interacting with the resulting selfie panorama 721. For instance, the user's blinks can be recorded and captured. Navigation of the selfie panorama can be done in a manner similar to that described in conjunction with FIG. 7J. In particular, the location and viewpoint of the person(s) in the selfie panorama 721 can be changed by the user by swiping/navigating in the direction of the main panorama axis. According to various embodiments, selfie panoramas 721 can also include parallax effects. These parallax effects can be seen when swiping/navigating perpendicular to the direction of the main panorama axis. In addition, similar to FIG. 7F, three-dimensional effects can be presented when swiping perpendicular to the panorama direction. In particular, when swiping perpendicular to the panorama direction, along the parallax direction, nearby objects are displaced along the parallax direction while the scene at distance stays still or moves less than the nearby objects.

As described above, various types of panoramas can be created with surround views. In addition, surround views can be viewed and navigated in different ways. With reference to FIG. 7M, shown is one example of extended views of panoramas that are provided based on user navigation. In the present example, possible views 727 include a full panorama view 729, recording views 731, and extended view 733. A full panorama view 729 includes a full view of the information in a surround view. The recording views 731 include the visual data captured in images and/or recordings. The extended view 733 shows more than what is visible during one point in time in recording views 731 but less than the full panorama view 729. The portion of the panorama 729 that is visible in an extended view 733 is defined by user navigation. An extended view 733 is especially interesting for a selfie or object panorama, because the extended view follows the object/person in the panorama and shows a larger view than what was visible for the camera while recording. Essentially, more context is provided to the user in an extended view 733 during navigation of the surround view.

According to various embodiments, once a series of images is captured, these images can be used to generate a surround view. With reference to FIG. 8, shown is an example of a surround view in which three-dimensional content is blended with a two-dimensional panoramic context. In the present example embodiment, the movement of capture device 820 follows a locally convex motion, such that the capture device moves around the object of interest (i.e., a person sitting in a chair). The object of interest is delineated as the content 808, and the surrounding scenery (i.e., the room) is delineated as the context 810. In the present embodiment, as the movement of the capture device 820 moves leftwards around the content 808, the direction of content rotation relative to the capture device 812 is in a rightward, counterclockwise direction. Views 802, 804, and 806 show a progression of the rotation of the person sitting in a chair relative to the room.

According to various embodiments, a series of images used to generate a surround view can be captured by a user recording a scene, object of interest, etc. Additionally, in some examples, multiple users can contribute to acquiring a series of images used to generate a surround view. With reference to FIG. 9, shown is one example of a space-time surround view being simultaneously recorded by independent observers.

In the present example embodiment, cameras 904, 906, 908, 910, 912, and 914 are positioned at different locations. In some examples, these cameras 904, 906, 908, 910, 912, and 914 can be associated with independent observers. For instance, the independent observers could be audience members at a concert, show, event, etc. In other examples, cameras 904, 906, 908, 910, 912, and 914 could be placed on tripods, stands, etc. In the present embodiment, the cameras 904, 906, 908, 910, 912, and 914 are used to capture views 904a, 906a, 908a, 910a, 912a, and 914a, respectively, of an object of interest 900, with world 902 providing the background scenery. The images captured by cameras 904, 906, 908, 910, 912, and 914 can be aggregated and used together in a single surround view in some examples. Each of the cameras 904, 906, 908, 910, 912, and 914 provides a different vantage point relative to the object of interest 900, so aggregating the images from these different locations provides information about different viewing angles of the object of interest 900. In addition, cameras 904, 906, 908, 910, 912, and 914 can provide a series of images from their respective locations over a span of time, such that the surround view generated from these series of images can include temporal information and can also indicate movement over time.

As described above with regard to various embodiments, surround views can be associated with a variety of capture modes. In addition, a surround view can include different capture modes or different capture motions in the same surround view. Accordingly, surround views can be separated into smaller parts in some examples. With reference to FIG. 10, shown is one example of separation of a complex surround-view into smaller, linear parts. In the present example, complex surround view 1000 includes a capture area 1026 that follows a sweeping L motion, which includes two separate linear motions 1022 and 1024 of camera 1010. The surround views associated with these separate linear motions can be broken down into linear surround view 1002 and linear surround view 1004. It should be noted that although linear motions 1022 and 1024 can be captured sequentially and continuously in some embodiments, these linear motions 1022 and 1024 can also be captured in separate sessions in other embodiments.

In the present example embodiment, linear surround view **1002** and linear surround view **1004** can be processed independently, and joined with a transition **1006** to provide a continuous experience for the user. Breaking down motion into smaller linear components in this manner can provide various advantages. For instance, breaking down these smaller linear components into discrete, loadable parts can aid in compression of the data for bandwidth purposes. Similarly, non-linear surround views can also be separated into discrete components. In some examples, surround views can be broken down based on local capture motion. For example, a complex motion may be broken down into a locally convex portion and a linear portion. In another example, a complex motion can be broken down into separate locally convex portions. It should be recognized that any number of motions can be included in a complex surround view **1000**, and that a complex surround view **1000** can be broken down into any number of separate portions, depending on the application.

Although in some applications, it is desirable to separate complex surround views, in other applications it is desirable to combine multiple surround views. With reference to FIG. **11**, shown is one example of a graph that includes multiple surround views combined into a multi-surround view **1100**. In this example, the rectangles represent various surround views **1102**, **1104**, **1106**, **1108**, **1110**, **1112**, **1114**, and **1116**, and the length of each rectangle indicates the dominant motion of each surround view. Lines between the surround views indicate possible transitions **1118**, **1120**, **1122**, **1124**, **1126**, **1128**, **1130**, and **1132** between them.

In some examples, a surround view can provide a way to partition a scene both spatially and temporally in a very efficient manner. For very large scale scenes, multi-surround view **1100** data can be used. In particular, a multi-surround view **1100** can include a collection of surround views that are connected together in a spatial graph. The individual surround views can be collected by a single source, such as a single user, or by multiple sources, such as multiple users. In addition, the individual surround views can be captured in sequence, in parallel, or totally uncorrelated at different times. However, in order to connect the individual surround views, there must be some overlap of content, context, or location, or of a combination of these features. Accordingly, any two surround views would need to have some overlap in content, context, and/or location to provide a portion of a multi-surround view **1100**. Individual surround views can be linked to one another through this overlap and stitched together to form a multi-surround view **1100**. According to various examples, any combination of capture devices with either front, back, or front and back cameras can be used.

In some embodiments, multi-surround views **1100** can be generalized to more fully capture entire environments. Much like “photo tours” collect photographs into a graph of discrete, spatially-neighboring components, multiple surround views can be combined into an entire scene graph. In some examples, this can be achieved using information obtained from but not limited to: image matching/tracking, depth matching/tracking, IMU, user input, and/or GPS. Within such a graph or multi-surround view, a user can switch between different surround views either at the end points of the recorded motion or wherever there is an overlap with other surround views in the graph. One advantage of multi-surround views over “photo tours” is that a user can navigate the surround views as desired and much more visual information can be stored in surround views. In contrast, traditional “photo tours” typically have limited views that can be shown to the viewer either automatically

or by allowing the user to pan through a panorama with a computer mouse or keystrokes.

According to various embodiments, a surround view is generated from a set of images. These images can be captured by a user intending to produce a surround view or retrieved from storage, depending on the application. Because a surround view is not limited or restricted with respect to a certain amount of visibility, it can provide significantly more visual information about different views of an object or scene. More specifically, although a single viewpoint may be ambiguous to adequately describe a three-dimensional object, multiple views of the object can provide more specific and detailed information. These multiple views can provide enough information to allow a visual search query to yield more accurate search results. Because a surround view provides views from many sides of an object, distinctive views that are appropriate for search can be selected from the surround view or requested from a user if a distinctive view is not available. For instance, if the data captured or otherwise provided is not sufficient to allow recognition or generation of the object or scene of interest with a sufficiently high certainty, a capturing system can guide a user to continue moving the capturing device or provide additional image data. In particular embodiments, if a surround view is determined to need additional views to produce a more accurate model, a user may be prompted to provide additional images.

With reference to FIG. **12**, shown is one example of a process for prompting a user for additional images **1200** to provide a more accurate surround view. In the present example, images are received from a capturing device or storage at **1202**. Next, a determination is made whether the images provided are sufficient to allow recognition of an object of interest at **1204**. If the images are not sufficient to allow recognition of an object of interest, then a prompt is given for the user to provide additional image(s) from different viewing angles at **1206**. In some examples, prompting a user to provide one or more additional images from different viewing angles can include suggesting one or more particular viewing angles. If the user is actively capturing images, the user can be prompted when a distinct viewing angle is detected in some instances. According to various embodiments, suggestions to provide one or more particular viewing angles can be determined based on the locations associated with the images already received. In addition, prompting a user to provide one or more additional images from different viewing angles can include suggesting using a particular capture mode such as a locally concave surround view, a locally convex surround view, or a locally flat surround view, depending on the application.

Next, the system receives these additional image(s) from the user at **1208**. Once the additional images are received, a determination is made again whether the images are sufficient to allow recognition of an object of interest. This process continues until a determination is made that the images are sufficient to allow recognition of an object of interest. In some embodiments, the process can end at this point and a surround view can be generated.

Optionally, once a determination is made that the images are sufficient to allow recognition of an object of interest, then a determination can then be made whether the images are sufficient to distinguish the object of interest from similar but non-matching items at **1210**. This determination can be helpful especially when using visual search, examples of which are described in more detail below with regards to FIGS. **19-22**. In particular, an object of interest may have distinguishing features that can be seen from particular

angles that require additional views. For instance, a portrait of a person may not sufficiently show the person's hairstyle if only pictures are taken from the front angles. Additional pictures of the back of the person may need to be provided to determine whether the person has short hair or just a pulled-back hairstyle. In another example, a picture of a person wearing a shirt might warrant additional prompting if it is plain on one side and additional views would show prints or other insignia on the sleeves or back, etc.

In some examples, determining that the images are not sufficient to distinguish the object of interest from similar but non-matching items includes determining that the number of matching search results exceeds a predetermined threshold. In particular, if a large number of search results are found, then it can be determined that additional views may be needed to narrow the search criteria. For instance, if a search of a mug yields a large number of matches, such as more than 20, then additional views of the mug may be needed to prune the search results.

If the images are not sufficient to distinguish the object of interest from similar but non-matching items at **1210**, then a prompt is given for the user to provide additional image(s) from different viewing angles at **1212**. In some examples, prompting a user to provide one or more additional images from different viewing angles can include suggesting one or more particular viewing angles. If the user is actively capturing images, the user can be prompted when a distinct viewing angle is detected in some instances. According to various embodiments, suggestions to provide one or more particular viewing angles can be determined based on the locations associated with the images already received. In addition, prompting a user to provide one or more additional images from different viewing angles can include suggesting using a particular capture mode such as a locally concave surround view, a locally convex surround view, or a locally flat surround view, depending on the application.

Next, the system receives these additional image(s) from the user at **1214**. Once the additional images are received, a determination is made again whether the images are sufficient to distinguish the object of interest from similar but non-matching items. This process continues until a determination is made that the images are sufficient to distinguish the object of interest from similar but non-matching items. Next, the process ends and a surround view can be generated from the images.

With reference to FIGS. **13A-13B**, shown are examples of prompts requesting additional images from a user in order to produce a more accurate surround view. In particular, a device **1300** is shown with a search screen. In FIG. **13A**, an example of a visual search query **1302** is provided. This visual search query **1302** includes an image of a white mug. The results **1306** include various mugs with a white background. In particular embodiments, if a large amount of search results is found, a prompt **1304** can be provided to request additional image data from the user for the search query.

In FIG. **13B**, an example of another visual search query **1310** is provided in response to prompt **1304** in FIG. **13A**. This visual search query **1310** provides a different viewpoint of the object and provides more specific information about the graphics on the mug. This visual search query **1310** yields new results **1312** that are more targeted and accurate. In some examples, an additional prompt **1308** can be provided to notify the user that the search is complete.

Once a surround view is generated, it can be used in various applications, in particular embodiments. One application for a surround view includes allowing a user to

navigate a surround view or otherwise interact with it. According to various embodiments, a surround view is designed to simulate the feeling of being physically present in a scene as the user interacts with the surround view. This experience depends not only on the viewing angle of the camera, but on the type of surround view that is being viewed. Although a surround view does not need to have a specific fixed geometry overall, different types of geometries can be represented over a local segment of a surround view such as a concave, convex, and flat surround view, in particular embodiments.

In particular example embodiments, the mode of navigation is informed by the type of geometry represented in a surround view. For instance, with concave surround views, the act of rotating a device (such as a smartphone, etc.) can mimic that of rotating a stationary observer who is looking out at a surrounding scene. In some applications, swiping the screen in one direction can cause the view to rotate in the opposite direction. This effect is akin to having a user stand inside a hollow cylinder and pushing its walls to rotate around the user. In other examples with convex surround views, rotating the device can cause the view to orbit in the direction it is leaning into, such that the object of interest remains centered. In some applications, swiping the screen in one direction causes the viewing angle to rotate in the same direction: this creates the sensation of rotating the object of interest about its axis or having the user rotate around the object. In some examples with flat views, rotating or moving a device can cause the view to translate in the direction of the device's movement. In addition, swiping the screen in one direction can cause the view to translate in the opposite direction, as if pushing foreground objects to the side.

In some examples, a user may be able to navigate a multi-surround view or a graph of surround views in which individual surround views can be loaded piece by piece and further surround views may be loaded when necessary (e.g. when they are adjacent to/overlap the current surround view and/or the user navigates towards them). If the user reaches a point in a surround view where two or more surround views overlap, the user can select which of those overlapping surround views to follow. In some instances, the selection of which surround view to follow can be based on the direction the user swipes or moves the device.

With reference to FIG. **14**, shown is one example of a process for navigating a surround view **1400**. In the present example, a request is received from a user to view an object of interest in a surround view at **1402**. According to various embodiments, the request can also be a generic request to view a surround view without a particular object of interest, such as when viewing a landscape or panoramic view. Next, a three-dimensional model of the object is accessed at **1404**. This three-dimensional model can include all or a portion of a stored surround view. For instance, the three-dimensional model can be a segmented content view in some applications. An initial image is then sent from a first viewpoint to an output device at **1406**. This first viewpoint serves as a starting point for viewing the surround view on the output device.

In the present embodiment, a user action is then received to view the object of interest from a second viewpoint. This user action can include moving (e.g. tilting, translating, rotating, etc.) an input device, swiping the screen, etc., depending on the application. For instance, the user action can correspond to motion associated with a locally concave surround view, a locally convex surround view, or a locally flat surround view, etc. According to various embodiments,

an object view can be rotated about an axis by rotating a device about the same axis. For example, the object view can be rotated along a vertical axis by rotating the device about the vertical axis. Based on the characteristics of the user action, the three-dimensional model is processed at **1410**. For instance, movement of the input device can be detected and a corresponding viewpoint of the object of interest can be found. Depending on the application, the input device and output device can both be included in a mobile device, etc. In some examples, the requested image corresponds to an image captured prior to generation of the surround view. In other examples the requested image is generated based on the three-dimensional model (e.g. by interpolation, etc.). An image from this viewpoint can be sent to the output device at **1412**. In some embodiments, the selected image can be provided to the output device along with a degree of certainty as to the accuracy of the selected image. For instance, when interpolation algorithms are used to generate an image from a particular viewpoint, the degree of certainty can vary and may be provided to a user in some applications. In other examples, a message can be provided to the output device indicating if there is insufficient information in the surround view to provide the requested images.

In some embodiments, intermediate images can be sent between the initial image at **1406** and the requested image at **1412**. In particular, these intermediate images can correspond to viewpoints located between a first viewpoint associated with the initial image and a second viewpoint associated with the requested image. Furthermore, these intermediate images can be selected based on the characteristics of the user action. For instance, the intermediate images can follow the path of movement of the input device associated with the user action, such that the intermediate images provide a visual navigation of the object of interest.

With reference to FIG. **15**, shown is an example of swipe-based navigation of a surround view. In the present example, three views of device **1500** are shown as a user navigates a surround view. In particular, the input **1510** is a swipe by the user on the screen of device **1500**. As the user swipes from right to left, the object of interest moves relative to the direction of swipe **1508**. Specifically, as shown by the progression of images **1506**, **1504**, and **1502**, the input **1510** allows the user to rotate around the object of interest (i.e., the man wearing sunglasses).

In the present example, a swipe on a device screen can correspond to rotation of a virtual view. However, other input modes can be used in other example embodiments. For instance, a surround view can also be navigated by tilting a device in various directions and using the device orientation direction to guide the navigation in the surround view. In another example, the navigation can also be based on movement of the screen by the user. Accordingly, a sweeping motion can allow the user to see around the surround view as if the viewer were pointing the device at the object of interest. In yet another example, a website can be used to provide interaction with the surround view in a web-browser. In this example, swipe and/or motion sensors may be unavailable, and can be replaced by interaction with a mouse or other cursor or input device.

According to various embodiments, surround views can also include tagging that can be viewed during navigation. Tagging can provide identification for objects, people, products, or other items within a surround view. In particular, tagging in a surround view is a very powerful tool for presenting products to users/customers and promoting those elements or items. In one example, a tag **1512** can follow the location of the item that is tagged, such that the item can be

viewed from different angles while the tag locations still stay valid. The tags **1512** can store different types of data, such as a name (e.g. user name, product name, etc.), a description, a link to a website/webshop, price information, a direct option for purchasing a tagged object, a list of similar objects, etc. In some examples, the tags can become visible when a user selects an item in a surround view. In other examples, the tags can be automatically displayed. In addition, additional information can be accessed by selecting a tag **1512** in some applications. For instance, when a user selects a tag, additional information can be displayed on screen such as a description, link, etc.

In some embodiments, a user can create a tag **1512** by selecting either a point or a region in one viewpoint of a surround view. This point or region is then automatically propagated into other viewpoints. Alternatively, tag locations can be automatically suggested to the user by an application based on different information, such as face detection, object detection, objects in focus, objects that are identified as foreground, etc. In some examples, object detection can be made from a database of known objects or object types/classes.

In the present example, tag **1512** identifies a shirt in the surround view. Of course, any text or title can be included, such as a name, brand, etc. This tag **1512** can be mapped to a particular location in the surround view such that the tag is associated with the same location or point in any view selected. As described above, tag **1512** can include additional information that can be accessed by tapping or otherwise selecting the tag, in some embodiments. Although tagging is shown in FIG. **15**, it should be noted that surround views may not include tagging in some examples.

According to various embodiments, surround views can be stored and accessed in various ways. In addition, surround views can be used in many applications. With reference to FIG. **16A**, shown are examples of a sharing service for surround views on a mobile device **1602** and browser **1604**. The mobile device **1602** and browser **1604** are shown as alternate thumbnail displays **1600**, because the surround views can be accessed by either interface, depending on the application. According to various embodiments, a set of surround views can be presented to a user in different ways, including but not limited to: a gallery, a feed, and/or a website. For instance, a gallery can be used to present a collection of thumbnails to a user. These thumbnails can be selected from the surround views either by the user or automatically. In some examples, the size of the thumbnails can vary based on characteristics such as, but not limited to: an automatically selected size that is based on the structure and size of the content it contains; and/or the popularity of the surround view. In another example, a feed can be used to present surround views using interactive thumbnails.

In the present example, surround view thumbnails from a mobile device **1602** include thumbnails **1604** and title/label/description **1604**. The thumbnails **1604** can include an image from the surround view. The title/label/description **1604** can include information about the surround view such as title, file name, description of the content, labels, tags, etc.

Furthermore, in the present example, surround view thumbnails from a browser **1604** include thumbnails **1606**, title/label/description **1608**, and notifications **1610**. The thumbnails **1606** can include an image from the surround view. The title/label/description **1608** can include information about the surround view such as title, file name, description of the content, labels, tags, etc. In addition, notifications **1610** can include information such as comments on a surround view, updates about matching content,

suggested content, etc. Although not shown on the mobile version, notifications can also be included, but may be omitted in the interest of layout and space considerations in some embodiments. In some examples, notifications can be provided as part of a surround view application on a mobile device.

With reference to FIG. 16B, shown are examples of surround view-related notifications on a mobile device. In particular, alternative notification screens 1620 for a device 1622 are shown that include different formats for notifications. In some examples, a user can navigate between these screens depending on the user's preferences.

In the present example, screen 1624 includes a notification 1626 that includes a recommendation to the user based on content from recent surround views. In particular, the recommendation relates to a trip to Greece based on the application's finding that the user has an affinity for statues. This finding can be inferred from content found in the user's stored or recently browsed surround views, in some examples.

In the present example, screen 1628 includes notifications 1630 based on content from surround views that the user has stored, browsed, etc. For instance, one notification is a recommendation for a pair of shoes available at a nearby retailer that are similar to the user's shoes as provided in a surround view model. The recommendation also includes a link to a map to the retailer. This recommendation can be based on a surround view that the user has saved of a pair of shoes. The other notification is a recommendation to connect to another user that shares a common interest/hobby. In this example, the recommendation is based on the user's detected interest in hats. These recommendations can be provided automatically in some applications as "push" notifications. The content of the recommendations can be based on the user's surround views or browsing history, and visual search algorithms, such as those described with regard to FIGS. 19-22, can be used in some examples.

Screen 1630 shows another form of notification 1632 in the present example. Various icons for different applications are featured on screen 1630. The icon for the surround view application includes a notification 1632 embedded into the icon that shows how many notifications are waiting for the user. When the user selects the icon, the notifications can be displayed and/or the application can be launched, according to various embodiments.

According to various embodiments of the present disclosure, surround views can be used to segment, or separate, objects from static or dynamic scenes. Because surround views include distinctive 3D modeling characteristics and information derived from image data, surround views provide a unique opportunity for segmentation. In some examples, by treating an object of interest as the surround view content, and expressing the remaining of the scene as the context, the object can be segmented out and treated as a separate entity. Additionally, the surround view context can be used to refine the segmentation process in some instances. In various embodiments, the content can be chosen either automatically or semi-automatically using user guided interaction. One important use for surround view object segmentation is in the context of product showcases in e-commerce, an example of which is shown in FIG. 17B. In addition, surround view-based object segmentation can be used to generate object models that are suited for training artificial intelligence search algorithms that can operate on large databases, in the context of visual search applications.

With reference to FIG. 17, shown is one example of a process for providing object segmentation 1700. At 1702, a first surround view of an object is obtained. Next, content is selected from the first surround view at 1704. In some examples, the content is selected automatically without user input. In other examples, the content is selected semi-automatically using user-guided interaction. The content is then segmented from the first surround view at 1706. In some examples, the content is segmented by reconstructing a model of the content in three-dimensions based on the information provided in the first surround view, including images from multiple camera viewpoints. In particular example embodiments, a mechanism for selecting and initializing a segmentation algorithm based on iterative optimization algorithms (such as graphical models) can be efficiently employed by reconstructing the object of interest, or parts of it, in three-dimensions from multiple camera viewpoints available in a surround view. This process can be repeated over multiple frames, and optimized until segmentation reaches a desired quality output. In addition, segmenting the content can include using the context to determine parameters of the content.

In the present example, once the content is segmented from the first surround view, a second surround view is generated that includes the object without the content or scenery surrounding the object. At 1708, this second surround view is provided. In some examples, the second surround view can then be stored in a database. This second surround view can be used in various applications. For instance, the segmented content includes a product for use in e-commerce. As illustrated in FIG. 17B, the segmented content can be used to show a product from various viewpoints. Another application includes using the second surround view as an object model for artificial intelligence training. In yet another application, the second surround view can be used in 3D printing. In this application, data from the second surround view is to a 3D printer.

Although the present example describes segmenting out content from a first surround view, it should be noted that context can also be segmented out in other examples. For instance, the background scenery can be segmented out and presented as a second surround view in some applications. In particular, the context can be selected from the first surround view and the context can be segmented from the first surround view, such that the context is separated into a distinct interactive model. The resulting surround view would then include the scenery surrounding an object but exclude the object itself. A segmented context model can also be used in various applications. For instance, data from the resulting surround view can be sent to a 3D printer. In some examples, this could be printed as a panoramic background on a flat or curved surface. If a context model is also printed, then the object of interest can be placed in front of the panoramic background to produce a three-dimensional "photograph" or model of the surround view. In another application, the segmented out context can be used as background to a different object of interest. Alternatively, a segmented out content can be placed in a new segmented out context. In these examples, providing an alternative content or context allows objects of interest to be placed into new backgrounds, etc. For instance, a surround view of a person could be placed in various background contexts, showing the person standing on a beach in one surround view, and standing in the snow in another surround view.

With reference to FIG. 17B, shown is one example of a segmented object viewed from different angles. In particular, a rotational view 1720 is shown of an athletic shoe.

Object views **1722**, **1724**, **1726**, **1728**, and **1730** show the athletic shoe from various angles or viewpoints. As shown, the object itself is shown without any background or context. According to various embodiments, these different views of the segmented object can be automatically obtained from surround view content. One application of these types of rotational views is in e-commerce to show product views from different angles. Another application can be in visual search, according to various embodiments.

According to various embodiments, surround views can be generated from data obtained from various sources and can be used in numerous applications. With reference to FIG. **18**, shown is a block diagram illustrating one example of various sources that can be used for surround view generation and various applications that can be used with a surround view. In the present example, surround view generation and applications **1800** includes sources for image data **1808** such as internet galleries **1802**, repositories **1804**, and users **1806**. In particular, the repositories can include databases, hard drives, storage devices, etc. In addition, users **1806** can include images and information obtained directly from users such as during image capture on a smartphone, etc. Although these particular examples of data sources are indicated, data can be obtained from other sources as well. This information can be gathered as image data **1808** to generate a surround view **1810**, in particular embodiments.

In the present example, a surround view **1810** can be used in various applications. As shown, a surround view can be used in applications such as e-commerce **1812**, visual search **1814**, 3D printing **1816**, file sharing **1818**, user interaction **1820**, and entertainment **1822**. Of course, this list is only illustrative, and surround views can also be used in other applications not explicitly noted.

As described above with regard to segmentation, surround views can be used in e-commerce **1812**. For instance, surround views can be used to allow shoppers to view a product from various angles. In some applications, shoppers can even use surround views to determine sizing, dimensions, and fit. In particular, a shopper can provide a self-model and determine from surround views whether the product would fit the model. Surround views can also be used in visual search **1814** as described in more detail below with regard to FIGS. **19-22**. Some of the visual search applications can also relate to e-commerce, such as when a user is trying to find a particular product that matches a visual search query.

Another application of segmentation includes three-dimensional printing (3D printing) **1816**. Three-dimensional printing has been recently identified as one of the future disruptive technologies that will improve the global economy in the next decade. According to various embodiments, content can be 3D printed from a surround view. In addition, the panoramic background context in a surround view can also be printed. In some examples, a printed background context can complement the final 3D printed product for users that would like to preserve memories in a 3D printed format. For instance, the context could be printed either as a flat plane sitting behind the 3D content, or as any other geometric shape (spherical, cylindrical, U shape, etc).

As described above with regard to FIG. **16A**, surround views can be stored with thumbnail views for user access. This type of application can be used for file sharing **1818** between users in some examples. For instance, a site can include infrastructure for users to share surround views in a

manner similar to current photo sharing sites. File sharing **1818** can also be implemented directly between users in some applications.

Also as described with regard to FIGS. **14** and **15**, user interaction is another application of surround views. In particular, a user can navigate through a surround view for their own pleasure or entertainment. Extending this concept to entertainment **1822**, surround views can be used in numerous ways. For instance, surround views can be used in advertisements, videos, etc.

As previously described, one application of surround views is visual search. FIGS. **19**, **20**, and **22** depict examples of visual search using surround views. According to various embodiments, using surround views can provide much higher discriminative power in search results than any other digital media representation to date. In particular, the ability to separate content and context in a surround view is an important aspect that can be used in visual search.

Existing digital media formats such as 2D images are unsuitable for indexing, in the sense that they do not have enough discriminative information available natively. As a result, many billions of dollars are spent in research on algorithms and mechanisms for extracting such information from them. This has resulted in satisfactory results for some problems, such as facial recognition, but in general the problem of figuring out a 3D shape from a single image is ill-posed in existing technologies. Although the level of false positives and negatives can be reduced by using sequences of images or 2D videos, the 3D spatial reconstruction methods previously available are still inadequate.

According to various embodiments, additional data sources such as location-based information, which are used to generate surround views, provide valuable information that improves the capability of visual recognition and search. In particular example embodiments, two components of a surround view, the context and the content, both contribute significantly in the visual recognition process. In particular example embodiments, the availability of three-dimensional information that the content offers can significantly reduce the number of hypotheses that must be evaluated to recognize a query object or part of a scene. According to various embodiments, the content's three-dimensional information can help with categorization (i.e., figuring out the general category that an object belongs to), and the two-dimensional texture information can indicate more about a specific instance of the object. In many cases, the context information in a surround view can also aid in the categorization of a query object, by explaining the type of scene in which the query object is located.

In addition to providing information that can be used to find a specific instance of an object, surround views are also natively suited for answering questions such as: "what other objects are similar in shape and appearance?" Similar to the top-N best matches provided in response to a web search query, a surround view can be used with object categorization and recognition algorithms to indicate the "closest matches," in various examples.

Visual search using surround views can be used and/or implemented in various ways. In one example, visual search using surround views can be used in object recognition for robotics. In another example, visual search using surround views can be used in social media curation. In particular, by analyzing the surround view data being posted to various social networks, and recognizing objects and parts of scenes, better #hashtags indices can be automatically generated. By generating this type of information, feeds can be curated and the search experience can be enhanced.

Another example in which visual search using surround views can be used is in a shopping context that can be referred to as “Search and Shop.” In particular, this visual search can allow recognition of items that are similar in shape and appearance, but might be sold at different prices in other stores nearby. For instance, with reference to FIG. 21, a visual search query may yield similar products available for purchase.

In yet another example in which visual search using surround views can be used is in a shopping context that can be referred to as “Search and Fit.” According to various embodiments, because surround view content is three-dimensional, precise measurements can be extracted and this information can be used to determine whether a particular object represented in a surround view would fit in a certain context (e.g., a shoe fitting a foot, a lamp fitting a room, etc).

In another instance, visual search using surround views can also be used to provide better marketing recommendation engines. For example, by analyzing the types of objects that appear in surround views generated by various users, questions such as “what type of products do people really use in their daily lives” can be answered in a natural, private, and non-intrusive way. Gathering this type of information can facilitate improved recommendation engines, decrease and/or stop unwanted spam or marketing ads, thereby increasing the quality of life of most users. FIG. 16B shows one implementation in which recommendations can be provided according to various embodiments of the present invention.

With reference to FIG. 19, shown is one example of a process for providing visual search of an object 1900, where the search query includes a surround view of the object and the data searched includes three-dimensional models. At 1902, a visual search query that includes a first surround view is received. This first surround view is then compared to stored surround views at 1904. In some embodiments, this comparison can include extracting first measurement information for the object in the first surround view and comparing it to second measurement information extracted from the one or more stored surround views. For instance, this type of measurement information can be used for searching items such as clothing, shoes, or accessories.

Next, a determination is made whether any stored surround views correspond to the first surround view at 1906. In some examples, this determination is based on whether the subject matter in any of the stored surround views is similar in shape to the object in the first surround view. In other examples, this determination is based on whether any of the subject matter in the stored surround views is similar in appearance to the object in the first surround view. In yet other examples, this determination is based on whether any subject matter in the stored surround views include similar textures included in the first surround view. In some instances, this determination is based on whether any of the contexts associated with the stored surround views match the context of the first surround view. In another example, this determination is based on whether the measurement information associated with a stored surround view dimensionally fits the object associated with the first surround view. Of course any of these bases can be used in conjunction with each other.

Once this determination is made, a ranked list of matching results is generated at 1908. In some embodiments, generating a ranked list of matching results includes indicating how closely any of the stored surround views dimensionally fits the object associated with the first measurement information. According to various embodiments, this ranked list

can include displaying thumbnails of matching results. In some examples, links to retailers can be included with the thumbnails. Additionally, information about the matching results such as name, brand, price, sources, etc. can be included in some applications.

Although the previous example includes using a surround view as a visual search query to search through stored surround views or three-dimensional models, current infrastructure still includes a vast store of two-dimensional images. For instance, the internet provides access to numerous two-dimensional images that are easily accessible. Accordingly, using a surround view to search through stored two-dimensional images for matches can provide a useful application of surround views with the current two-dimensional infrastructure.

With reference to FIG. 20, shown is one example of a process for providing visual search of an object 2000, where the search query includes a surround view of the object and the data searched includes two-dimensional images. At 2002, a visual search query that includes a first surround view is received. Next, object view(s) are selected from the surround view at 2004. In particular, one or more two-dimensional images are selected from the surround view. Because these object view(s) will be compared to two-dimensional stored images, selecting multiple views can increase the odds of finding a match. Furthermore, selecting one or more object views from the surround view can include selecting object views that provide recognition of distinctive characteristics of the object.

In the present example, the object view(s) are then compared to stored images at 2006. In some embodiments, one or more of the stored images can be extracted from stored surround views. These stored surround views can be retrieved from a database in some examples. In various examples, comparing the one or more object views to the stored images includes comparing the shape of the object in the surround view to the stored images. In other examples, comparing the one or more object views to the stored images includes comparing the appearance of the object in the surround view to the stored images. Furthermore, comparing the one or more object views to the stored images can include comparing the texture of the object in the surround view to the stored images. In some embodiments, comparing the one or more object views to the stored images includes comparing the context of the object in the surround view to the stored images. Of course any of these criteria for comparison can be used in conjunction with each other.

Next, a determination is made whether any stored images correspond to the object view(s) at 2008. Once this determination is made, a ranked list of matching results is generated at 2010. According to various embodiments, this ranked list can include displaying thumbnails of matching results. In some examples, links to retailers can be included with the thumbnails. Additionally, information about the matching results such as name, brand, price, sources, etc. can be included in some applications.

With reference to FIG. 21, shown is an example of a visual search process 2100. In the present example, images are obtained at 2102. These images can be captured by a user or pulled from stored files. Next, according to various embodiments, a surround view is generated based on the images. This surround view is then used as a visual search query that is submitted at 2104. In this example, a surround view can be used to answer questions such as “which other objects in a database look like the query object.” As illustrated, surround views can help shift the visual search paradigm from finding other “images that look like the

query,” to finding other “objects that look like the query,” due to their better semantic information capabilities. As described with regard to FIGS. 19 and 20 above, the surround view can then be compared to the stored surround views or images and a list of matching results can be provided at 2106.

Although the previous examples of visual search include using surround views as search queries, it may also be useful to provide search queries for two-dimensional images in some embodiments. With reference to FIG. 22, shown is an example of a process for providing visual search of an object 2200, where the search query includes a two-dimensional view of the object and the data searched includes surround view(s). At 2202, a visual search query that includes a two-dimensional view of an object to be searched is received. In some examples, the two-dimensional view is obtained from an object surround view, wherein the object surround view includes a three-dimensional model of the object. Next, the two-dimensional view is compared to surround views at 2204. In some examples, the two-dimensional view can be compared to one or more content views in the surround views. In particular, the two-dimensional view can be compared to one or more two-dimensional images extracted from the surround views from different viewing angles. According to various examples, the two-dimensional images extracted from the surround views correspond to viewing angles that provide recognition of distinctive characteristics of the content. In other examples, comparing the two-dimensional view to one or more surround views includes comparing the two-dimensional view to one or more content models. Various criteria can be used to compare the images or models such as the shape, appearance, texture, and context of the object. Of course any of these criteria for comparison can be used in conjunction with each other.

With reference to FIG. 23, shown is a particular example of a computer system that can be used to implement particular examples of the present invention. For instance, the computer system 2300 can be used to provide surround views according to various embodiments described above. According to particular example embodiments, a system 2300 suitable for implementing particular embodiments of the present invention includes a processor 2301, a memory 2303, an interface 2311, and a bus 2315 (e.g., a PCI bus). The interface 2311 may include separate input and output interfaces, or may be a unified interface supporting both operations. When acting under the control of appropriate software or firmware, the processor 2301 is responsible for such tasks such as optimization. Various specially configured devices can also be used in place of a processor 2301 or in addition to processor 2301. The complete implementation can also be done in custom hardware. The interface 2311 is typically configured to send and receive data packets or data segments over a network. Particular examples of interfaces the device supports include Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like.

In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management.

According to particular example embodiments, the system 2300 uses memory 2303 to store data and program instructions and maintained a local side cache. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to store received metadata and batch requested metadata.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to tangible, machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include hard disks, floppy disks, magnetic tape, optical media such as CD-ROM disks and DVDs; magneto-optical media such as optical disks, and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and programmable read-only memory devices (PROMs). Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

Although many of the components and processes are described above in the singular for convenience, it will be appreciated by one of skill in the art that multiple components and repeated processes can also be used to practice the techniques of the present disclosure.

While the present disclosure has been particularly shown and described with reference to specific embodiments thereof, it will be understood by those skilled in the art that changes in the form and details of the disclosed embodiments may be made without departing from the spirit or scope of the invention. It is therefore intended that the invention be interpreted to include all variations and equivalents that fall within the true spirit and scope of the present invention.

What is claimed is:

1. A method comprising:

generating a combined surround view including a first surround view placed in a foreground position relative to a second surround view, wherein the first and second surround views are obtained using different capture motions, and wherein non-linear surround views are separated into discrete components and complex motions are broken down into locally convex and linear portions such that a more complex surround view corresponding to a more complex capture motion can be generated using separate smaller surround views corresponding to separate smaller capture motions.

2. The method of claim 1, wherein the first surround view is an object panorama.

3. The method of claim 1, wherein the first surround view includes a panoramic multi-view interactive digital media representation of an object.

4. The method of claim 1, wherein the second surround view includes a panoramic view of a distant scene.

5. The method of claim 1, wherein the combined surround view includes dynamic content that is animated during user navigation.

6. The method of claim 1, wherein an object in the combined surround view appears three-dimensional when the combined surround view is navigated.

7. The method of claim 1, wherein an object is displayed on a device, wherein a viewing angle of the object is manipulated by tilting and rotating the device.

31

- 8.** A system comprising:
a processor; and
memory, wherein the memory includes instructions for:
generating a combined surround view including a first
surround view placed in a foreground position relative
to a second surround view, wherein the first and second
surround views are obtained using different capture
motions, and wherein non-linear surround views are
separated into discrete components and complex
motions are broken down into locally convex and linear
portions such that a more complex surround view
corresponding to a more complex capture motion can
be generated using separate smaller surround views
corresponding to separate smaller capture motions.
- 9.** The system of claim **8**, wherein the first surround view
is an object panorama.
- 10.** The system of claim **8**, wherein the first surround view
includes a panoramic multi-view interactive digital media
representation of an object.
- 11.** The system of claim **8**, wherein the second surround
view includes a panoramic view of a distant scene.
- 12.** The system of claim **8**, wherein the combined sur-
round view includes dynamic content that is animated
during user navigation.
- 13.** The system of claim **8**, wherein an object in the
combined surround view appears three-dimensional when
the combined surround view is navigated.
- 14.** The system of claim **8**, wherein an object is displayed
on a device, wherein a viewing angle of the object is
manipulated by tilting and rotating the device.

32

- 15.** A non-transitory computer readable medium compris-
ing instructions for:
generating a combined surround view including a first
surround view placed in a foreground position relative
to a second surround view, wherein the first and second
surround views are obtained using different capture
motions, and wherein non-linear surround views are
separated into discrete components and complex
motions are broken down into locally convex and linear
portions such that a more complex surround view
corresponding to a more complex capture motion can
be generated using separate smaller surround views
corresponding to separate smaller capture motions.
- 16.** The non-transitory computer readable medium of
claim **15**, wherein the first surround view is an object
panorama.
- 17.** The non-transitory computer readable medium of
claim **15**, wherein the first surround view includes a pan-
oramic multi-view interactive digital media representation
of an object.
- 18.** The non-transitory computer readable medium of
claim **15**, wherein the second surround view includes a
panoramic view of a distant scene.
- 19.** The non-transitory computer readable medium of
claim **15**, wherein the combined surround view includes
dynamic content that is animated during user navigation.
- 20.** The non-transitory computer readable medium of
claim **15**, wherein an object in the combined surround view
appears three-dimensional when the combined surround
view is navigated.

* * * * *